



OneFS

OneFS 6.5.5 User Guide

Contents

Introduction to Isilon scale-out NAS solutions.....	9
OneFS architecture overview.....	9
File system scalability.....	10
Cluster communications overview.....	10
Cluster splitting and merging.....	10
Quorums.....	11
File system journals.....	11
Network management.....	12
Internal network management.....	12
Internal network settings.....	12
Modify the netmask for the internal network.....	13
Add IP addresses to the internal network.....	13
Delete IP addresses from the internal network.....	14
Migrate internal network IP addresses.....	14
Enable the int-b/failover network	15
Modify the int-b/failover network.....	15
External network management.....	16
DNS management.....	16
External network settings.....	18
IP address pool configuration.....	25
SmartConnect management	27
NIC aggregation.....	33
VLAN management.....	36
Node provisioning rules.....	37
Cluster management.....	39
Cluster monitoring.....	39
View cluster status.....	39
View node status.....	40
Monitor cluster size.....	40
Monitor cluster throughput.....	41
Monitor CPU usage.....	41
View client connections.....	41
View cluster and node statistics.....	42
View active alerts.....	42
View cluster logs.....	42
SNMP monitoring.....	42
Configure SNMP monitoring.....	43
Download MIBs.....	44
Cluster configuration.....	45
Configure the OneFS web Interface port.....	45

Add nodes to a cluster.....	45
Remove nodes from a cluster.....	46
Upgrade cluster operating system.....	46
Shut down or reboot the cluster.....	46
Set the cluster name.....	47
Specify company contact information.....	47
Specify email server settings.....	47
Manually set the cluster date and time.....	48
Set the cluster join mode.....	48
Cluster services.....	49
Configure telnet.....	49
Configure NTP.....	49
Configure character set encoding.....	49
Enable access time tracking.....	50
SAS drive LED status.....	50
File system management.....	51
File sharing.....	51
Multi-protocol file sharing.....	52
SMB (Windows file sharing protocol).....	52
NFS (UNIX file sharing protocol).....	58
HTTP and DAV.....	62
FTP.....	63
OneFS data protection.....	64
Data layout and file striping.....	64
Isilon FlexProtect.....	65
N+M data protection.....	65
Data mirroring.....	66
Metadata and inodes.....	66
Protection level management.....	67
Data backup.....	67
Smart failure and recovery.....	67
Drive failures.....	68
Node failures.....	68
WORM (write once, read many).....	68
File System Explorer.....	72
Navigate the OneFS file structure.....	73
View file and directory properties.....	73
Modify protection settings.....	74
Modify I/O optimization settings.....	75
Modify UNIX permissions.....	76
Add a directory.....	77
NDMP backup management.....	77
Configure NDMP backup settings.....	77
Manage NDMP backup devices.....	78

Manage NDMP backup ports.....	79
View NDMP backup logs.....	80
View and manage NDMP backup sessions.....	80
Cluster anti-virus scanning.....	81
Cluster anti-virus summary.....	81
The cluster anti-virus scanning service.....	82
ICAP scan server configuration.....	82
Anti-virus global settings.....	85
Anti-virus scanning policies.....	86
Anti-virus threat responses.....	90
Anti-virus scan reports.....	92
Authentication, identity management, and authorization.....	94
Authentication sources.....	94
View authentication service status.....	94
Local users and groups.....	95
File provider.....	98
Active Directory.....	99
LDAP.....	101
Legacy LDAP.....	103
NIS.....	105
Advanced authentication settings.....	107
Configure general authentication settings.....	107
Run the Repair Permissions job.....	108
Identity management.....	108
Access tokens.....	109
User mapping.....	110
ID mapping.....	112
On-disk identity selection.....	113
Authorization.....	115
File authorization data.....	115
ACLs and UNIX-style permissions.....	115
Authorization data mapping in mixed environments.....	120
Job management.....	121
Monitor jobs.....	121
Modify a job.....	121
Start a job.....	122
Control a job.....	123
Update a job.....	123
View job history.....	123
Impact policy management.....	124
View impact policy settings.....	124
Create an impact policy.....	124
Copy an impact policy.....	125
Modify an impact policy.....	125

Delete an impact policy.....	125
SmartPools.....	126
SmartPools overview.....	126
SmartPools monitoring.....	126
SmartPools configuration.....	127
Configure basic SmartPools settings.....	127
Configure default protection settings.....	128
Configure default I/O optimization settings.....	129
Reprovision the cluster.....	129
Disk pool management.....	130
Monitor disk pools.....	130
Create a disk pool.....	130
Modify a disk pool.....	131
Delete a disk pool.....	131
File pool policy management.....	131
View file pool policies.....	132
Add a file pool policy.....	132
Modify a file pool policy.....	134
Prioritize a file pool policy.....	134
Copy a file pool policy.....	135
Remove a file pool policy.....	135
Use a file pool policy template.....	135
Modify the default file pool policy.....	135
Software module licensing.....	137
View module license status information.....	137
Activate a module license through the web interface.....	137
Activate a module license through the command-line interface.....	138
Unconfigure a module license.....	138
The Isilon SyncIQ module.....	139
SyncIQ snapshot overview.....	139
SyncIQ policies and jobs.....	140
SyncIQ policy configuration.....	140
SyncIQ policy management.....	149
SyncIQ policy and job operations.....	150
SyncIQ performance.....	154
View SyncIQ performance statistics.....	154
View SyncIQ performance rules.....	154
Create a SyncIQ network-usage rule.....	155
Create a SyncIQ file-operation rule.....	155
Modify a SyncIQ performance rule.....	156
Delete a SyncIQ performance rule.....	156
SyncIQ reports.....	156
View SyncIQ reports.....	157
Configure global default SyncIQ report settings.....	157

SynclQ events.....	158
View and interpret SynclQ events.....	158
Manage SynclQ event alerts.....	158
The Isilon SmartQuotas module.....	159
SmartQuotas overview.....	159
SmartQuotas upgrades.....	161
Upgrade to SmartQuotas 2.0.....	161
Quota management.....	162
Create an accounting quota.....	162
Create an enforcement quota.....	162
Search for a quota.....	163
View quota settings.....	164
Modify a quota.....	165
Clone a quota.....	165
Delete a quota.....	165
Unlink a quota from a default quota.....	166
Import a SmartQuotas 2.0 configuration file.....	166
Import a SmartQuotas 1.x configuration file.....	166
Export a quota configuration file.....	167
Quota notifications.....	167
Configure default global notification settings.....	168
Configure basic custom quota notification settings.....	168
Configure advanced custom quota notification settings.....	169
Map an email notification for a quota.....	171
Configure a custom email notification template.....	172
Quota reports.....	173
Configure quota report settings.....	173
View quota reports.....	173
Run a live quota report.....	174
Download a quota report.....	175
The Isilon SnapshotIQ module.....	176
Snapshot schedule settings.....	176
Snapshot schedule configuration.....	176
Snapshot schedule management.....	179
Manual snapshots.....	180
Manually create a snapshot.....	180
Snapshot management.....	180
View SnapshotIQ summary information.....	180
View a recent snapshot.....	181
View a pending snapshot.....	181
Rename or modify a snapshot.....	181
Delete a snapshot.....	182
Snapshot usage and reserve settings.....	182
View individual snapshot usage.....	182

View overall snapshot usage.....	183
Configure snapshot reserve levels.....	183
File and folder restoration.....	184
Restore a deleted file using Shadow Copy Emulation.....	184
Restore a corrupted or overwritten file using Shadow Copy Emulation.....	184
Restore a deleted folder using Shadow Copy Emulation.....	185
SnapshotIQ settings.....	185
Enable or disable SnapshotIQ.....	185
Configure basic SnapshotIQ settings.....	185
Configure advanced SnapshotIQ settings.....	186
The Isilon iSCSI module.....	188
iSCSI overview.....	188
Security.....	188
Supported iSCSI initiators.....	189
Limitations and considerations.....	189
Global iSCSI settings.....	189
Monitor iSCSI sessions.....	189
Configure the iSCSI service.....	190
Configure the iSNS client service.....	190
Configure default SmartConnect pools.....	191
iSCSI target management.....	191
Create a target.....	191
View target settings.....	192
Modify a target.....	193
Delete a target.....	193
Initiator access control.....	194
Configure access control settings.....	194
Add an initiator to the access list.....	194
Modify initiator settings.....	195
Remove an initiator from the access list.....	195
CHAP authentication.....	195
Enable or disable CHAP authentication.....	195
Create a CHAP secret.....	196
Modify a CHAP secret.....	196
Delete a CHAP secret.....	197
iSCSI LUN management.....	197
Create a logical unit.....	197
View logical unit settings.....	199
Control access to a logical unit.....	200
Modify a logical unit.....	200
Delete a logical unit.....	200
Clone a logical unit.....	200
Move a logical unit.....	201
Import a logical unit.....	202

Apache Hadoop.....	203
Isilon and Hadoop cluster integration.....	203
Create a local Hadoop user.....	203
Configure a Hadoop client.....	204
Configure HDFS.....	204
Enable or disable the HDFS service.....	206
The Isilon SupportIQ module.....	207
Data collected by SupportIQ.....	207
SupportIQ scripts.....	208
SupportIQ configuration and management.....	210
Enable and configure SupportIQ.....	210
Disable SupportIQ.....	211
Events and event notifications.....	212
View events and event notifications.....	212
View the event summary.....	212
View the event history.....	212
View event notification rules.....	212
View events settings.....	213
View event help.....	213
Event management.....	213
Quiet an event.....	213
Send a test event.....	213
Event notification management.....	214
Add an event notification rule.....	214
Edit an event notification rule.....	214
Event notification reference.....	214
System disk events.....	214
Node status events.....	218
Reboot events.....	224
Software events.....	226
SmartQuotas events.....	231
Snapshot events.....	232
Windows Networking events.....	235
File System events.....	240
Hardware events.....	241

Introduction to Isilon scale-out NAS solutions

The Isilon suite of scale-out NAS systems and software enables you to store and manage digital content and other unstructured data. An Isilon cluster consists of three or more nodes, which form a single unified file system.

The Isilon OneFS distributed file system, developed by Isilon Systems, is the patent-pending software that runs on Isilon's enterprise-class hardware and provides the intelligence behind all Isilon clustered storage systems.

OneFS provides the following features:

- Fully distributed single file system
- High-performance, fully symmetric cluster
- File striping across nodes
- Automated software to eliminate complexity
- Dynamic content balancing
- Flexible data protection
- High availability
- Web-based and command-line administration

Isilon OneFS enables independent and linear scalability of performance and capacity. OneFS provides a single point of management for large and rapidly growing repositories of unstructured and file-based data and offers fast access to large files, inherent high availability, and the ability to easily scale a cluster's throughput and capacity as your storage needs change.

Each node in an Isilon clustered storage system is a peer, so any node can handle a request. Using InfiniBand or Gigabit Ethernet for intracluster communication and synchronization, OneFS provides each node with knowledge of the entire file system layout.

Isilon OneFS controls the placement of files directly on individual disks. By laying out information on disks in a file-by-file manner, OneFS can control the redundancy level of the storage system at the volume, directory, and file levels.

OneFS architecture overview

Each Isilon cluster creates a single namespace and file system. There is no partitioning, and no need for volume creation. Because all information is shared among nodes, the entire file system is accessible by clients connecting to any node in the cluster.

All nodes in an Isilon cluster are fully symmetric. Because all nodes in the cluster are peers, the Isilon clustered storage system does not have any "master" or "slave" nodes. All data is striped across all nodes in the cluster. Metadata is also distributed across the cluster, and every node has access to the data in the event of a failure.

As nodes are added, the file system grows dynamically and content is redistributed. Each Isilon storage node contains globally coherent RAM, meaning that, as a cluster becomes larger, it also becomes faster. Each time a node is added, the cluster's concurrent performance scales linearly.

The minimum size of any Isilon cluster is three nodes, while the required minimum protection level depends on the type of Isilon node. For most Isilon node types, the minimum cluster size of three nodes ensures the minimum protection level of N+1. However, for 4U Isilon IQ X-Series and NL-Series nodes, and IQ 12000X/EX 12000 combination platforms, the minimum cluster size of three nodes requires a minimum protection level of N+2:1.

All Isilon IQ clusters use high-speed, low-latency InfiniBand fabric for intracluster communication. A cluster can grow to a maximum of 144 nodes. Clusters are expected to be composed of nodes of like models, with some exceptions made

for node types that are no longer available or when new models are introduced. For example, a cluster composed of IQ 6000X nodes can be expanded by adding IQ 9000X nodes.

You can add Isilon EX storage-expansion nodes to an existing Isilon IQ cluster using a SAS connection to increase storage space, in effect doubling the disk capacity of the node without increasing CPU or RAM. Isilon EX 6000 expansion nodes are supported only when used with Isilon IQ 6000X storage nodes, and the same applies to EX 9000 expansion nodes paired only with IQ 9000X nodes, and EX 12000 expansion nodes paired only with IQ 12000X nodes. In addition, there is a one-to-one relationship between IQ and EX nodes: For example, to use the EX 6000, each IQ 6000X node in the cluster must have an EX 6000 attached at the same time.

You can also add Isilon performance-accelerator nodes to an existing cluster to increase overall throughput. Acting as a new node, the Isilon IQ Accelerator node adds CPU, RAM, and an additional network connection to a cluster without adding any additional storage disks. Typically, if you are supporting write-intensive applications, one Accelerator node per storage node is recommended. For read-intensive applications, up to three Accelerator nodes per storage node are recommended.

The Isilon IQ Accelerator-x node provides 10-gigabit Ethernet connections for high-single-stream client connections. The Accelerator-X also includes more CPU and memory resources to support a greater client load.

The Isilon IQ Backup Accelerator node provides enhanced high-speed backup of file-system data to locally attached tape or media-changer devices that are connected through Fibre Channel on the back end. The Backup Accelerator node supports a wide range of data management applications (DMAs), tape libraries, and switches.

File system scalability

Adding a new node requires no downtime. Scaling a cluster requires no reconfiguration, no server or client mount points, and no application changes.

OneFS can scale to provide multiple petabytes of storage in a single file system, so there is no need to create small volumes or logical units. As the cluster scales, Isilon AutoBalance migrates content to new storage nodes while the system is online and in production. Data is automatically balanced across all nodes, reducing costs, complexity, and risk. In addition, as nodes are added to the cluster, Isilon TrueScale linearly scales system throughput.

Cluster communications overview

Isilon clusters use separate internal and external networks.

The internal network, also called the back-end network, uses Gigabit Ethernet (GigE) or InfiniBand, with the option to configure an additional failover network for redundancy. Essentially, the back-end network acts as the backplane of the cluster, enabling each node to act as a contributor to the whole. It is recommended that you avoid using the back-end network for any other purpose.

Clients connect to the cluster using external (front-end) GigE connections. The Isilon cluster supports standard network communication protocols, including UNIX file sharing (NFS), Windows file sharing (SMB), HTTP, and FTP. The cluster includes multiple front-end Ethernet connections, providing flexibility for a wide variety of network configurations.

Cluster splitting and merging

The cluster as a whole monitors every node, enabling any one node to know the status of any other at any time. If a node is not reachable over the cluster interconnect, it is “split” from the cluster, acting as though the cluster were split into two separate groups of nodes.

While in a split state, all data is reachable and modifiable on the side of the majority of nodes. If the node becomes reachable again, a “merge” occurs, bringing that node back into the cluster. (The two groups merge back into one.) The

node can rejoin the cluster without being rebuilt and reconfigured. This is unlike RAID arrays, which require drives to be rebuilt.

Quorums

The Isilon system uses a quorum to prevent “split-brain” conditions that can be introduced if the cluster should temporarily split into two clusters.

In order for the cluster to properly function and accept data writes, a quorum of nodes must be active and responding. A quorum is defined as more than half of the nodes in the cluster. For example, in a seven-node cluster, a four-node quorum would be required. In a 10-node cluster, six nodes would be required. If a node or group of nodes is operational and responsive, but is not a member of a quorum, it runs in a read-only state.

File system journals

OneFS is a journaled file system in which each node contains a battery-backed NVRAM card that is used for journaling. The NVRAM card battery charge lasts up to three days without system power.

A file system journal, which stores information about changes to the file system before they are written to the disk, is designed to enable fast, consistent recoveries after system failures or crashes, such as power loss. When a node or cluster boots up, the file system checks its journal and replays any entries that it finds.

OneFS will mount only if it can guarantee that all transactions not already in the system have been recorded. For example, if proper shutdown procedures were not followed, and the NVRAM battery discharged, transactions might have been lost; to prevent any potential problems, the node will not mount the file system. If this situation occurs, contact Isilon Technical Support for assistance.

Network management

After you determine the topology of your network, you can set up and manage your internal and external networks.

Two types of networks are associated with a cluster:

- **Internal:** Communication occurs through InfiniBand connections. Optionally, you can configure an additional failover network for redundancy. The back-end network acts as the backplane of the cluster, enabling each node to act as a contributor to the whole. It is recommended that you avoid using the back-end internal network for any purpose other than intra-cluster communication.
- **External:** Clients connect to the cluster through the external Ethernet connections. The Isilon cluster supports standard network communication protocols, including UNIX file sharing (NFS), Windows file sharing (SMB), HTTP, and FTP. The cluster includes various front-end Ethernet connections, providing flexibility for a wide variety of network configurations. Front-end speeds vary across various products.

The cluster's web administration interface enables you to manage both the internal and external network settings from a centralized location.

Internal network management

Internal networks enable communication between Isilon nodes in a cluster. You can configure a single internal network, or, optionally, you can specify a second internal network that includes internal network failover if either the int-a or int-b port fails.

To enable an internal failover network, the int-a ports of each node in the cluster must be physically connected to one switch, and the int-b ports on each node must be physically connected to another switch. When the three following conditions are met, the failover function is automatically enabled:

- Both internal networks are present.
- IP Ranges for both the int-b and Failover internal networks are configured.
- Both the int-a and int-b interfaces are enabled.



Important: If your network topology will use a single internal network, only the int-a interface must be configured. If your network topology will use more than one internal network, you must first configure the int-a interfaces, and then configure the int-b and failover interfaces.

Internal network settings

To configure the initial settings for the cluster's internal networks(int-a, int-b, and failover), you must use command-line Configuration Wizard.

To configure the int-a network settings, follow the prompts in the command-line wizard to perform the following actions:

- Set the netmask value
- Add IP addresses

After you have configured the int-a network settings, you can use the web administration interface to:

- Delete IP addresses
- Migrate IP addresses
- Configure initial settings for the int-b and failover networks
- Modify the int-b and failover network settings

- Enable internal network failover (optional)

Modify the netmask for the internal network

You can modify the netmask value for the internal network.

Netmask settings require modification if the currently defined netmask is too restrictive for the size of the internal network.



Important: To enable changes to the netmask value, you must reboot the cluster.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **Internal Network Settings**, click **int-a** or **int-b / Failover**.
The **Edit Internal Network** page for the selected network appears. The **Settings** area shows the current **Netmask** value.
3. In the **Netmask** box, type a netmask value.



Note: You cannot modify the netmask value if the change would cause any nodes' addresses to become invalid.

4. Click **Submit**.
A dialog box appears and prompts you to reboot the cluster.
5. Specify when you want to reboot the cluster:
 - To immediately reboot the cluster, click **Yes**. A rebooting-status page appears while the cluster reboots. When the rebooting process is complete, the login page appears.
 - Click **No** to return to the **Edit Internal Network** page without changing the settings or rebooting the cluster.

When the rebooting process is complete, the login page appears.

Add IP addresses to the internal network

You can add IP addresses to the internal network. Having a sufficient number of IP addresses ensures that new nodes can be joined to the cluster.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **Internal Network Settings**, click **int-a** or **int-b / Failover**.
The **Edit Internal Network** page for the selected network appears. The **IP Ranges** area displays the currently configured IP addresses for the internal network.



Note: The **IP Ranges** section includes separate areas for migrating int-b and failover networks.

3. Click **Add range**.
The **Add IP Range** dialog box appears.
4. In the first **IP range** box, type the IP address at the low end of the range.
5. In the second **IP range** box, type the IP address at the high end of the range.
6. Click **Submit**.
The **Edit Internal Network** page appears, and the new IP range appears in the **IP Ranges** list.



Note: Adding addresses does not affect the addresses of nodes currently in your cluster, and does not require a reboot in order to take effect.

Delete IP addresses from the internal network

You can remove IP addresses from the internal network to decrease the IP address range.

Example cases include preparing to migrate IP address ranges, preparing to remove nodes from the cluster, or reconfiguring the internal network.



Important: You cannot delete IP addresses that are currently in use by a node on the cluster.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **Internal Network Settings**, click either **int-a** or **int-b / Failover**.
The **Edit Internal Network** page for the selected network appears. The **IP Ranges** area displays the currently configured IP addresses for the internal network.



Note: The **IP Ranges** section includes separate areas for migrating int-b and failover networks.

3. Delete IP address ranges as needed:
 - To delete a portion of an IP address range, click **Delete range**. The **Delete IP Range** dialog box appears.
 - In the first **IP range** box, type the IP address at the low end of the range.
 - In the second **IP range** box, type the IP address at the high end of the range.
 - Click **Submit**.
 - To delete a specific IP address range, in the **IP Ranges** list, click **Delete** next to the IP address range. The **Confirm** dialog box appears. Click **Yes**.
4. Click **Submit**.
The **Edit Internal Network** page refreshes.



Note: Deleting IP addresses does not affect the addresses of nodes currently in the cluster, and does not require you to reboot the cluster in order to take effect. However, before you *change* any IP addresses that are in use by nodes in the cluster, you must first migrate the addresses.

Migrate internal network IP addresses

You can migrate IP addresses for your internal network in order to change addresses that are currently in use by nodes in the cluster.



Important: If you migrate addresses that are in use by more than one node, you must reboot the cluster to ensure that all processes are using the correct IP address.

You can migrate IP addresses on your internal network in order to change addresses that are currently in use by nodes in the cluster.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **Internal Network Settings**, click **int-a** or **int-b / Failover**.
The **Edit Internal Network** page appears for the selected network.
The **IP Ranges** section displays the network's currently configured IP addresses.
3. Migrate IP ranges as needed:
 - To migrate only a portion of an IP address range, click **Migrate range**. The **Migrate IP Range** dialog box appears.

In the **Current IP range** and **New IP range** boxes, type the low and high IP addresses, and then click **Submit**. The **Edit Internal Network** page appears.

- To migrate a specific range in the **IP Ranges** list, click **Migrate** next to the range's name. The **Migrate IP Range** dialog box appears, displaying the current IP address range in the **Current IP range** boxes.

In the **New IP range** boxes, type the low and high IP addresses, and then click **Submit**. The **Edit Internal Network** page appears.

4. If necessary, modify the netmask value in the **Netmask** box.
5. Click **Submit**.
The **Confirm Cluster Reboot** dialog box appears.
6. Click **Yes** to reboot the cluster. A rebooting-status page appears while the cluster reboots. When the reboot process is complete, the cluster login page appears.

Enable the int-b/failover network

You can enable the int-b and failover internal networks to provide back-up networks in the event of an int-a network failure. By default, the int-b and failover internal networks are disabled.



Note: You must configure the int-b and failover internal network settings before you can enable them. Configuration involves specifying a valid netmask and IP address range for the network.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **Internal Network Settings**, click **int-b / Failover**.
The **Edit Internal Network** page for the int-b and failover networks appears.
3. Click **Enable** next to the **State** option to enable int-b and failover networks.



Important: To enable the changes to the network settings, you must reboot the cluster.

4. Specify a valid netmask ensuring that there is no overlap between int-a, int-b, and failover.



Note: Specifying a netmask value is required.

5. Click **Submit**.
The **Confirm Cluster Reboot** dialog box appears.
6. To reboot the cluster, click **Yes**.
The cluster-reboot process begins. When the reboot process is complete, the login page appears.

Modify the int-b/failover network

You can modify the int-b and failover network settings to ensure that the int-b and failover IP address ranges can accommodate all nodes in the cluster.

This section describes how to modify these settings through the web administration interface.



Important: To enable changes to the netmask, you must reboot the cluster.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **Internal Network Settings**, click **int-b / Failover**.
The **Edit Internal Network** page for the int-b and failover networks appears.
3. Modify the network's netmask settings as needed:

- In the **State** area, click to **Disable** or **Enable** the internal network.
 - To modify the netmask setting, in the **Netmask** box, type a new netmask value.
4. Click **Submit**.
The **Confirm Cluster Reboot** dialog box appears.
 5. To reboot the cluster, click **Yes**.
The cluster-reboot process begins. When the reboot process is complete, the login page appears.

External network management

External networks provide communication outside of the cluster. OneFS supports network subnets, IP address pools, and network provisioning rules to facilitate the external-network configuration process.

Subnets simplify external network management, and provide flexibility in implementing and maintaining efficient cluster network strategies. By creating IP address pools within subnets, you can further partition your network interfaces. Using provisioning rules, external network settings can be configured once then automatically applied as nodes are added to a cluster.



Note: You must initially configure the default external IP subnet in IPv4 format. After the initial configuration is complete, you can configure additional subnets using IPv4 or IPv6.

IP address pools can be associated with a node or a group of nodes, and with specific NIC ports on the nodes. For example, you might decide to establish one subnet for storage nodes, and a different profile for accelerator nodes, based on expected network traffic volume.

How you specify your external network subnets will depend on the topology of your network. In a basic network topology where each node communicates to the clients on the same subnet, only a single external subnet is required. However, in a more complex topology where some nodes connect to one external IP subnet, additional nodes connect to a second IP subnet, and other nodes do not connect externally at all, several different external network subnets will be required.

DNS management

You can configure the DNS settings you use for your external network.

To address the cluster using friendly names instead of IP addresses, you must name them on the DNS server, and make the DNS server accessible to the cluster.

The DNS property designates one or two Domain Name Service (DNS) devices the cluster uses to resolve friendly names to IP addresses. If you designate two devices, you must also designate one as the primary DNS device and the other as the secondary device. Isilon IQ will automatically failover to the secondary DNS device if the primary DNS device goes offline.

You can configure the DNS settings during initial cluster configuration using the command-line Configuration Wizard, or at any time after initial setup using the web administration interface, or the `isi networks` command via the command-line-interface.

Configure the DNS settings

You can configure the domain name servers and DNS search list that the cluster uses to resolve host names.

You may have already configured DNS settings for the cluster during initial configuration of your Isilon cluster using the command-line Configuration Wizard, but you can also use the web administration interface to configure or modify the domain name server(s) and DNS search list the cluster uses to resolve host names.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. In the **DNS Settings** area, click **Edit**.
The **Configure DNS** dialog box appears.

3. In the **Domain name server(s)** box, type the IP address of the name server you want to add.

This is the domain name server address that the cluster uses to answer all DNS requests.



Note: You can specify domain name server addresses in IPv4 or IPv6 format.

4. In the **DNS search list** box, enter the local domain name.



Note: The domain name you type in the **DNS search list** is used for resolving unqualified hostnames.

5. Click **Submit**.

The **Networking** page appears and displays the new name server record and DNS search list under **DNS Settings**.

Add a name server record to the DNS infrastructure

You can add a new name server record to an existing DNS configuration using the web administration interface.

Prerequisite: SmartConnect requires that a new name server (NS) record be added to the existing authoritative DNS zone that contains the cluster. In the Microsoft Windows DNS Management utility, this type of record is called a New Delegation, which is just an NS record.



Note: This topic describes the process of adding a new NS record to an existing authoritative DNS zone, using the Microsoft Windows DNS Management utility. For the UNIX or Linux environment, a new NS record can be added using BIND.

1. On the Windows taskbar, click **Start**, point to **All Programs**, select **Administrative Tools**, and then click **DNS** to open the DNS Management utility.
2. In the left navigation pane of the DNS Management utility, expand the **Forward Lookup Zones** folder.
3. In the list of domain names in the **Forward Lookup Zones** folder, click to select your domain.
4. On the **Action** menu, click **New Delegation...** to open the **New Delegation** wizard.



Note: A delegated domain name is required only if you are configuring your cluster to use either SmartConnect Basic or SmartConnect Advanced for connection balancing.

5. On the **New Delegation wizard Welcome view**, click **Next** to open the **Delegated Domain Name** view.
6. In the **Delegated** domain text box, type an authoritative DNS domain name.



Note: The authoritative DNS domain name you type in the **Delegated** domain box is responsible for answering host name requests.

7. Click **Next** to open the **Name Servers** view.
8. Click **Add** to open the **New Resource Record** dialog box.
9. In the **Server fully qualified domain name (FQDN)** box, type the FQDN that appeared in Step 6.
10. In the **IP_address** box, type the virtual IP address.



Note: Note: The IP address that you type in this **IP_address** box is the virtual IP address used by SmartConnect to answer all DNS requests sent to its zone. In selecting a virtual IP address, the IP address must be on the same subnet it will be answering requests for, but not in the range of IP addresses reserved for nodes in the cluster.

11. Click **Add** to add the name server to the list.
12. Click **OK** to return to the **Name Servers** view.
13. Click **Next**.
14. Click **Finish**.

External network settings

You can configure initial settings for your external networks through a wizard in the command-line interface.

During initial cluster setup, you must specify the following information about your external network:

- Netmask
- IP address range
- Gateway
- Domain name server list (optional)
- DNS search list (optional)
- SmartConnect zone name (optional)
- SmartConnect service address (optional)

After you configure these initial settings, OneFS automatically performs the following actions:

- It creates a default external network subnet named *subnet0*, with the specified netmask, gateway and optional SmartConnect service address.
- It creates a default IP address pool named *pool0* with the specified IP address range, the optional SmartConnect zone name, and the initial external interface of the first node in the cluster as the only member.
- It creates a default network provisioning rule named *rule0*, which automatically assigns the first external interface for all newly added nodes to *pool0*.
- It adds *pool0* to *subnet0* and configures it to use the virtual IP of *subnet0* as its SmartConnect service address.
- It sets the global, outbound DNS settings to the optional domain name server list and DNS search list, if provided.

After you have configured the cluster's initial network configuration through the command-line Configuration Wizard, you can make additional changes to your external network settings through the web administration interface. For example, you can add additional external network subnets, or modify existing external network settings including subnets, IP address pools, and network provisioning rules.

External network configuration

OneFS provides a four-step wizard for adding new external network subnets to a cluster.

The subnet wizard guides you through the following steps any time you add a new external network configuration to your cluster via the web administration interface:

1. Configuring the subnet's basic and advanced settings.
2. Assigning one or more IP address pools used by the subnet.
3. Optionally configuring connection balancing for the IP address pools.
4. Assigning network interfaces to the subnet's IP address pools.

Once you add an external network subnet to a cluster using the subnet wizard, you can modify those settings at any time.

Add an external network subnet

You can add a subnet to a cluster's external network using the subnet wizard.

OneFS provides a four-step wizard that enables you to add and configure an external network subnet. This procedure explains how to start the subnet wizard and perform the first of the four steps.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **External Network Settings**, click **Add subnet**.
The Configure Network Subnet wizard starts, and the **Step 1 of 4 — Subnet Settings** dialog box appears.
3. In the **Basic** section, in the **Name** box, type a unique name for the subnet.
The name can be up to 32 alphanumeric characters long and can include underscores or hyphens, but not spaces or other punctuation.
4. In the **Description** box, type an optional descriptive comment about the subnet.

This value can contain up to 128 characters. A description is optional, but can be helpful when you are later managing your networks.

5. Specify the subnet's IP address format and configure an associated netmask or prefix-length setting:
 - For an IPv4 subnet, in the **IP family** list click **IPv4** and then, in the **Netmask** box, type a dotted decimal octet (x.x.x.x) that represents the subnet's mask.
 - For an IPv6 subnet, in the **IP family** list click **IPv6** and then, in the **Prefix length** box, type an integer (ranging from 1 to 128) that represents the network interface's prefix length.
6. In the **MTU** list, type or select the maximum size of the transmission units the cluster uses in network communication. Any numerical value is allowed, but might not be compatible with your network. Common settings are **1500** (standard frames) and **9000** (jumbo frames).
 Although OneFS supports both 1500 MTU and 9000 MTU or larger packet sizes, it is recommended that you configure switches for jumbo frames. Jumbo frames enable the Isilon cluster to more efficiently communicate to all storage nodes in the cluster.
7. In the **Gateway address** box, type the IP address of the gateway server device through which the cluster communicates with systems outside the subnet.
8. In the **Gateway priority** box, type an integer that represents the priority of the subnet's gateway in relation to other subnets' gateways for nodes that are assigned to more than one subnet.
 You can configure only one default gateway per node, but each subnet can have its own gateway. When a node belongs to more than one subnet, this option enables you to define the preferred default gateway. A value of **1** represents the highest priority, and **10** represents the lowest priority.
9. If you plan to use SmartConnect for connection balancing, in the **SmartConnect service IP** box, type the IP address that will receive all incoming DNS requests from outside the cluster. SmartConnect answers these DNS requests for each IP address pool according to the pool's client connection policy. You must have at least one subnet configured with a SmartConnect service IP in order to use connection balancing.
10. In the **Advanced** section, you can optionally enable **VLAN tagging** if you want to enable the cluster to participate in virtual networks.



Caution: Configuring a VLAN requires advanced knowledge of how to configure network switches. Consult your network switch documentation before configuring your cluster for a VLAN.

11. If you enable VLAN tagging, you must also type a **VLAN ID** that corresponds to the ID number for the VLAN set on the switch, with a value from **2** to **4094**.
12. In the **Hardware load balancing** box, you can optionally type the IP address for a hardware load balancing switch using Direct Server Return (DSR). This routes all client traffic to the cluster via the switch. The switch in turn determines which node will handle the client's traffic and then passes the traffic to that node.
13. Click **Next**.
 The **Step 2 of 4 — IP Address Pool Settings** dialog box appears.

This is the first of four steps required to configure an external network subnet. To save the network configuration, you must complete each of the remaining steps. For information on the next step, see [Configure an IP address pool](#) on page 19.

Configure an IP address pool

You must configure an IP address pool for a cluster's external network subnet.

Prerequisite: This is the second of four steps required to configure an external network subnet on a cluster. You must complete the steps on the previous subnet wizard page before you can perform these steps.

IP address pools enable you to partition your cluster's external network interfaces into groups, or pools, of unique IP address ranges within a specified subnet.



Note: If your cluster is running SmartConnect Basic for connection balancing, you can configure only one IP address pool per subnet. With the optional, licensed SmartConnect Advanced module, you can configure unlimited IP address pools per subnet.

1. In the **Step 2 of 4 — IP Address Pool Settings** dialog box, type a unique **Name** for the IP address pool. The name can be up to 32 alphanumeric characters long and can include underscores or hyphens, but no spaces or other punctuation.
2. Type an optional **Description** for the IP address pool. The description can contain up to 128 characters.
3. In the **IP range (low-high)** area, click **New**.
The subnet wizard adds an IP address range with default **Low IP** and **High IP** values.
4. Click to select the default **Low IP** value and then replace it with the starting IP address of the subnet's IP address pool.
5. Click to select the default **High IP** value and then replace it with the ending IP address of the subnet's IP address pool.
6. Optionally, add additional IP address ranges to the IP address pool by repeating steps 3 through 5 as needed.
7. Click **Next**.

The **Step 3 of 4 — SmartConnect Settings** dialog box appears.

This is the second of four steps required to configure an external network subnet. To save the network configuration, you must complete each of the remaining steps. For information on the next step, see [Configure SmartConnect settings](#) on page 20.

Configure SmartConnect settings

You can configure connection balancing for a cluster's external network subnet.

Prerequisite: This is the third of four steps required to configure an external network subnet on a cluster using the subnet wizard. You must complete the steps on the previous subnet wizard page in order to access this page.

SmartConnect is a client connection balancing management module that can optionally be configured as part of a cluster's external network subnet. The settings available on this wizard page depend on whether you have SmartConnect Basic (provided by default) or SmartConnect Advanced, which is an optional licensed module.

1. In the **Step 3 of 4 — SmartConnect Settings** dialog box, type a **Zone name** for the SmartConnect zone that this IP address pool represents. The zone name must be unique among the pools served by the SmartConnect service subnet specified in Step 3 below.
2. In the **Connection policy** list, select the type of connection balancing policy the IP address pool for this subnet will use. The policy determines how SmartConnect distributes incoming DNS requests across the members of an IP address pool.

- **Round Robin:** This method selects the next available node on a rotating basis. This is the default state (once SmartConnect is activated) if no other policy is selected.



Note: Round robin is the only connection policy available with SmartConnect Basic. The following client connection balancing policies are only available with SmartConnect Advanced.

- **Connection Count:** This method determines the number of open TCP connections on each available node to select which node the client will connect to.
- **Network Throughput:** This method uses the overall average throughput volume on each available node to optimize the cluster usage.
- **CPU Usage:** This method examines average CPU utilization on each node to optimize the cluster usage.

3. In the **SmartConnect service subnet** list, select the name of the external network subnet whose SmartConnect service will answer DNS requests on behalf of the IP address pool. A pool can have only one SmartConnect service answering DNS requests. If this option is left blank, the IP address pool it belongs to will be excluded when SmartConnect answers incoming DNS requests for the cluster.

If you have purchased the optional SmartConnect Advanced module, complete the following steps for the options in the **SmartConnect Advanced** section of this wizard page.

4. In the **IP allocation method** list, select the method by which IP addresses are assigned to the member interfaces for this IP address pool:

- **Static:** Select this IP allocation method to assign IP addresses when member interfaces are added to the IP pool. As members are added to the pool, this method allocates the next unused IP address from the pool to each new member. Once an IP address is allocated, the pool member keeps the address indefinitely unless:
 - the member interface is removed from the network pool
 - the member node is removed from the cluster
 - the member interface is moved to another IP address pool
- **Dynamic:** Select this IP allocation method to ensure that all IP addresses in the IP address pool are assigned to member interfaces, which allows clients to connect to any IP addresses in the pool and be guaranteed a response. If a node or an interface becomes unavailable, its IP addresses are automatically moved to other available member interfaces in the pool.

If you select the dynamic IP allocation method, you can further customize the dynamic behavior of the IP address pool in the following steps by specifying the SmartConnect **Rebalance policy** and the **IP failover policy in the next two steps**.

5. Select which type of SmartConnect **Rebalance policy** to use when IP addresses are redistributed when node interface members in an IP address pool become available after a period of unavailability:
 - **Automatic Failback:** This policy (the default) automatically redistributes IP addresses. The automatic rebalance is triggered by a change to:
 - the cluster membership
 - the cluster's external network configuration
 - a member network interface
 - **Manual Failback:** This policy does not redistribute IP addresses until you manually issue a rebalance command via the command line or the web administration interface.
6. The **IP failover policy**—also known as NFS failover—determines how to redistribute the IP addresses among remaining members of an IP address pool when one or more members are unavailable. In order to enable NFS failover, you must first set the **IP allocation method** to **Dynamic**, and then select which type of **IP failover policy** to use:
 - **Round Robin:** This method reassigns the IP address of a failed node interface member on a rotating basis.
 - **Connection Count:** This method determines the number of open TCP connections on each available node, and then attempts to proportionally redistribute the IP addresses from failed node interface members to the remaining members.
 - **Network Throughput:** This method evaluates the overall throughput volume, and then redistributes the IP addresses from failed node interface members to the remaining members based on optimizing this volume.
 - **CPU Usage:** This method examines average CPU utilization on each node, and then attempts to proportionally distribute the IP addresses from a failed node interface member to the remaining members across all nodes in the cluster.

7. Click **Next** to store the changes you made to this wizard page.
The **Step 4 of 4 — IP Address Pool members** dialog box appears.

This is the third of four steps required to configure an external network subnet. To save the network configuration, you must complete each of the remaining steps. For information on the next step, see [Select the interface members for an IP address pool](#) on page 21.

Select the interface members for an IP address pool

You can select which external network interfaces belong to the IP address pool that belongs to the external network subnet.

Prerequisite: This is the final of four steps required to configure an external network subnet on a cluster using the subnet wizard. The choices available on this wizard page depend on the types of external network interfaces of your Isilon cluster nodes.

1. In the **Step 4 of 4 — IP Address Pool Members** dialog box, select which **Available interfaces** on which nodes you want to assign to the current IP address pool, and then click the right-arrow button.

You can also drag and drop the selected interfaces between the **Available interfaces** table and the **Interfaces in current pool** table.

Selecting an available interface for a node that has a **Type** designated **Aggregation** bonds together the external interfaces for the selected node. For information about NIC aggregation, refer to [NIC aggregation](#) on page 33 in the OneFS user guide.



Note: In the case of aggregated links, choose the aggregation mode that corresponds to the switch settings from the **Aggregation mode** drop-down.



Caution: Configuring link aggregation requires advanced knowledge of how to configure network switches. Consult your network switch documentation before configuring your cluster for link aggregation.

2. When you have finished assigning external network interfaces to the IP address pool, click **Submit**.
The external subnet settings you configured using the subnet wizard are displayed on the **Edit Subnet** page.

You can change the subnet configuration settings at any time without going through the four-step wizard process. See [External subnet management](#) on page 22 for details.

External subnet management

Once you have configured an external network subnet using the subnet wizard, you can modify the subnet as your networking needs evolve. Before you modify an existing subnet, consider how the changes will affect your cluster configuration and performance.

When you initially configure an external network subnet, the subnet wizard requires you to go through a sequential four-step process. To modify a subnet, however, you can go directly to any of the subnet settings you want to change.

You can also modify your cluster's default external network subnet, which was configured during cluster installation using the command-line Configuration Wizard.

Modify the default subnet

After you initially set up your Isilon cluster using the command-line Configuration Wizard, you may chose to make changes to the default external network settings that were made by the wizard.

You may also want to configure additional external network settings that were not part of the initial configuration, such as connection balancing or NFS failover. You can do this by modifying the settings of the default subnet, which is named *subnet0*.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **External Network Settings**, click **subnet0**, which is the name of the cluster's default external network subnet.
The **Edit Subnet** page appears.
3. Make any required changes to the following **Basic** subnet settings:
 - **Description:** A descriptive comment about the subnet that can be up to 128 characters long.
 - **Netmask:** An IP address that specifies the IP mask for a network interface.
 - **MTU:** The maximum size of the transmission units the cluster uses in network communication. Any numerical value is allowed, but might not be compatible with your network. Common settings are **1500** (standard frames) and **9000** (jumbo frames).
 - **Gateway address:** The IP address of the gateway server device through which the cluster communicates with systems outside the subnet.
 - **Gateway priority:** The priority of the subnet's gateway in relation to other subnets' gateways for nodes that are assigned to more than one subnet. Only one default gateway can be configured on each Isilon node, but each subnet can have its own gateway. When a node belongs to more than one subnet, this option allows you to define the preferable default gateway. A value of **1** is the highest priority, with **10** being the lowest priority.
 - **SmartConnect service IP:** The IP address that will receive all incoming DNS requests from outside the cluster. SmartConnect answers these DNS requests for each IP address pool according to the pool's client connection

policy. You must have at least one subnet configured with a SmartConnect service IP in order to use connection balancing.

4. Make any required changes to the following **Advanced** subnet settings:

- **VLAN tagging:** Optionally, enabling Virtual LAN (VLAN) tagging allows a cluster to participate in multiple virtual networks. VLAN support provides security across subnets that is otherwise available only by buying additional network switches.



Note: Configuring a VLAN requires advanced knowledge of how to configure your network switches to enable this option on a cluster. Consult your network switch documentation before configuring your cluster for a VLAN.

- **VLAN ID:** If VLAN tagging is enabled, you must type a **VLAN ID** that corresponds to the ID number for the VLAN set on the switch, with a value from **2** to **4094**.
- **Hardware load balancing:** The IP address for a hardware load balancing switch, if used. This routes all client traffic to the cluster via the switch. The switch in turn determines which node will handle the client's traffic and passes the traffic to the node.

5. When you have finished modifying the external network subnet settings, click **Submit**.

The **Edit Subnet** page for the subnet you modified appears.

If you need to make additional changes to your cluster's external network settings click one of the following options on the **Edit Subnet** page:

- To modify basic subnet settings such as the netmask, gateway address, and SmartConnect service IP, and advanced settings for VLAN and hardware load balancing, click **Edit** for **subnet0**.
- To modify the settings of the subnet's IP address pool, in the **IP address pools** area, click **Edit** next to **Basic Settings**. If subnet0 has more than one IP address pool, click the plus sign icon next to the pool you want to modify to view its details.
- To modify the SmartConnect settings for the subnet's IP address pool, in the **SmartConnect settings** area, click **Edit**.
- To modify the network interfaces that belong to the subnet's IP address pool, in the **Pool members** area, click **Edit**.

Modify an external subnet

You can modify any external network subnet settings.

When you initially configure an external network subnet, the subnet wizard requires you to go through a sequential four-step process. To modify a subnet, however, you can go directly to any of the subnet settings you want to change.




Caution: Modifying an external network subnet that is in use could prevent access to the Isilon cluster and the web administration interface. OneFS will display a warning if deleting a subnet will terminate communication between the cluster and the web administration interface.

1. On the **Cluster** menu, click **Networking** .
The **Networking** page appears.
2. Under **External Network Settings**, click the link for the subnet that you want to modify.
The **Edit Subnet** page appears.
3. In the **Settings** area, click **Edit**.
The **Configure Subnet** dialog box appears.
4. Modify **Basic** subnet settings as needed:
 - **Description** : A descriptive comment about the subnet that can be up to 128 characters long.
 - **Netmask** : An IP address that specifies the IP mask for the network interface. This setting is available for only IPv4 subnets.
 - **Prefix length** : An integer (ranging from 1 to 128) that specifies the network portion of the IP address. This setting is available for only IPv6 subnets.

- **MTU** : The maximum size of the transmission units the cluster uses in network communication. Any numerical value is allowed, but might not be compatible with your network. Common settings are **1500** (standard frames) and **9000** (jumbo frames).
- **Gateway address** : The IP address of the gateway server device through which the cluster communicates with systems outside of the subnet.
- **Gateway priority** : The priority of the subnet's gateway in relation to other subnets' gateways for nodes that are assigned to more than one subnet. Only one default gateway can be configured on each Isilon node, but each subnet can have its own gateway. If a node belongs to more than one subnet, this option enables you to define the preferred default gateway. A value of **1** is the highest priority, with **10** being the lowest priority.
- **SmartConnect service IP** : The IP address that will receive incoming DNS requests from outside the cluster. SmartConnect answers these DNS requests for each IP address pool according to the pool's client connection policy. You must have at least one subnet configured with a SmartConnect service IP in order to use connection balancing.

5. Modify **Advanced** subnet settings as needed:

- **VLAN tagging** : Optionally, enable Virtual LAN (VLAN) tagging, which allows a cluster to participate in multiple virtual networks. VLAN support provides security across subnets that is otherwise available only by buying additional network switches.
 **Note:** Configuring a VLAN requires advanced knowledge of how to configure your network switches. Consult your network switch documentation before configuring your cluster for a VLAN. If you are not using a virtual LAN, leave the VLAN options disabled.
- **VLAN ID** : If you enabled VLAN tagging, type a **VLAN ID** that corresponds to the ID number for the VLAN set on the switch, with a value from **1** to **4094**.
- **Hardware load balancing IPs** : The IP address for a hardware load balancing switch, if used. This routes all client traffic to the cluster via the switch. The switch in turn determines which node will handle the client's traffic and passes the traffic to the node.

6. Click **Submit** .

Delete an external subnet

If your network evolves over time and you no longer need a specific external network subnet, you can delete it.



Caution: Deleting an external network subnet that is in use could prevent access to the Isilon IQ cluster and the web administration interface. OneFS will display a warning if deleting a subnet will terminate communication between the cluster and the web administration interface.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **External Network Settings**, click the name of the subnet you want to delete.
The **Edit Subnet** page appears for the subnet you specified.
3. Click **Delete subnet**.
OneFS asks you to confirm that you want to delete the subnet.
4. Click **Yes** to delete the subnet, or click **No** if you want to retain it.
If the subnet you are deleting is used to communicate with the cluster's web administration interface, the confirmation message will contain an additional warning.
The **Edit Subnet** page appears.

IP address pool configuration

IP address pools are logical network partitions of the nodes and external network interfaces that belong to a cluster. IP address pools are also used to configure SmartConnect zones and IP failover support for protocols such as NFS.

IP address pools, which belong to external network subnets, allow you to partition your cluster's network interfaces into groups. This lets you assign ranges of IP addresses to logical or functional groups in your organization. For example, you could dedicate one range of IP addresses on a subnet to your organization's sales team, and another range of them to your human resources department.

A subnet's IP address pool consists of one or more ranges of unique and contiguous IP addresses, and the cluster's external network interfaces that they use to communicate with clients.

A default IP address pool is configured during initial setup of your cluster using the command-line Configuration Wizard. You can modify that default IP address pool at any time, add or remove additional pools, and modify them as your networking needs evolve.

If you add more external network subnets to your cluster using the subnet wizard, part of that process involves specifying the IP address pool or pools that belong to the subnet.

Each IP address pool is configured to allocate IP addresses to the cluster's external network interfaces using either the static or dynamic method. With the static allocation method, IP addresses are assigned when member interfaces are added to the pool; not all IP addresses are guaranteed to be assigned using this method. With the dynamic allocation method, all IP addresses are distributed across the member interfaces.

Add interface members to an IP address pool

You can assign a node's external network interfaces to a subnet's IP address pool.

Node interface members are initially assigned to a subnet's IP address pool during cluster setup, and additional members can be added later using the subnet wizard. Once a subnet and its IP address pool have been created, you can assign or reassign node interface members to the IP address pool at any time.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **External Network Settings**, click the link for the subnet that contains the IP address pool to which you want to add interface members.
The **Edit Subnet** page appears.
3. In the **Pool members** area, click **Edit**.
The **Configure Pool Interface Members** dialog box appears.
4. Select which **Available interfaces** on which nodes you want to assign to the current IP address pool, and then click the right-arrow button.

You can also drag and drop the selected interfaces between the **Available interfaces** table and the **Interfaces in current pool** table.

Selecting an available interface for a node that has a **Type** designated **Aggregation** bonds together the external interfaces for the selected node. For information about NIC aggregation, refer to [NIC aggregation](#) on page 33 in the OneFS user guide.



Note: In the case of aggregated links, choose the aggregation mode that corresponds to the switch settings from the **Aggregation mode** drop-down.



Caution: Configuring link aggregation requires advanced knowledge of how to configure network switches. Consult your network switch documentation before configuring your cluster for link aggregation.

5. When you have finished assigning external network interfaces to the IP address pool, click **Submit**.
The interface member settings you configured are displayed on the **Edit Subnet** page.

Remove interface members from an IP address pool

If you need to remove a node's external network interface from a subnet's IP address pool, you can do that from the cluster's web administration interface.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **External Network Settings**, click the name of the subnet that contains the IP address pool from which you want to remove the interface member(s).
The **Edit Subnet** page appears.
3. In the **Pool members** area, click **Edit**.
The **Configure Pool Interface Members** dialog box appears.
4. To remove a node's interface from the IP address pool, click its name in the **Interfaces in current pool** column, and then click the left-arrow button.
You can also drag and drop node interfaces between the **Available interfaces** column and the **Interfaces in current pool** column.
5. When you have finished removing node interfaces from the IP address pool, click **Submit**.
Your current interface member settings are displayed on the **Edit Subnet** page.

Allocate IP addresses to accommodate new nodes

As your organization's need for storage increases, you can easily expand capacity by adding new nodes to your Isilon cluster.

Once the hardware installation is complete, you may need to allocate IP addresses for the new node or nodes on one of the cluster's existing external network subnets, and then add the node's external interfaces to the subnet's IP address pool.

You can also use network provisioning rules to automate the process of configuring the external network interfaces for new nodes when they are added to a cluster, although you may still need to allocate more IP addresses for the new nodes, depending on how many are already configured.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **External Network Settings**, click the name of the subnet that contains the IP address pool that you want to allocate more IP address to in order to accommodate the new nodes.
The **Edit Subnet** page appears.
3. In the **Basic settings** area under **IP Address Pools** area, click **Edit**.
The **Configure IP address pool settings** dialog box appears.
4. To add new IP addresses, you can either:
 - Click **New** to add a new range of IP addresses using the **Low IP** and **High IP** columns.
 - Extend an existing IP range by clicking the value for either **Low IP** or **High IP**, and then typing a new beginning or ending IP address that results in a larger range for the pool.
5. Click **Submit**.
The new or expanded range of IP addresses appears on the **Edit Subnet** page.
6. In the **Pool members** area under **IP Address Pools**, click **Edit**.
The **Configure Pool Interface Members** dialog box appears.
7. In the Available Interfaces table, select one or more interfaces for the newly added node, and then click the right-arrow button to move them into the **Interfaces in current pool** table.
8. Click **Submit** to assign the new node interfaces to the IP address pool.
The new node interface assignments and IP addresses appear on the **Edit Subnet** page.

Move nodes between IP address pools

You can move nodes between IP address pools in the event of a network reconfiguration or installation of a new network switch.

The process of moving nodes between IP address pools involves creating a new IP address pool, assigning it to the nodes so that they are temporarily servicing multiple subnets. Once testing confirms that the new IP address pool is working correctly, the old IP address pool can safely be deleted.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **External Network Settings**, in the **Subnets** area, click the name of the subnet that will receive the new IP addresses for the nodes.
The **Edit Subnet** page appears.
3. In the **IP Address Pools** area, click **Add pool**.
The **Create new pool for subnet** dialog box appears.
4. Type a unique **Name** for the new IP address pool. The name can be up to 32 alphanumeric characters long and can include underscores or hyphens, but no spaces or other punctuation.
5. Type an optional **Description** for the IP address pool; it can be up to 128 characters long.
6. In the **IP range (low-high)** area, click **New**.
The subnet wizard adds an IP address range with default **Low IP** and **High IP** values.



Note: Depending on the IP version of the subnet family, you can enter either IPv4 or IPv6 addresses.

7. Click to select the default **Low IP** value and then replace it with the starting IP address of the subnet's IP address pool.
8. Click to select the default **High IP** value and then replace it with the ending IP address of the subnet's IP address pool.
9. Click **Next**.
The **SmartConnect Settings** dialog box appears.



Note: You can optionally enable SmartConnect for the pool, or skip that step and click **Next**.

10. In the **Available interfaces** list, click to select the interface or interfaces that you want to assign to the new IP address pool, and then click the right-arrow button.
You can also drag and drop the selected interfaces between the **Available interfaces** list and the **Interfaces in current pool** list.
11. Click **Submit**.
The new pool and its IP addresses, and any interface member settings that you configured, appear on the **Edit Subnet** page.
12. Verify that the new IP address pool is working correctly.
13. On the **Edit Subnet** page, under **IP address pools**, click **Delete pool** for the old pool that you want to replace with the new pool you just created.
OneFS prompts you to confirm that you want to delete the IP address pool.
14. Click **Yes**.
The updated interface member settings appear on the **Edit Subnet** page.

SmartConnect management

SmartConnect is a client connection balancing management module that enables client connections to be balanced across all nodes within an Isilon IQ cluster or across selected nodes.

SmartConnect is available in two versions:

- The SmartConnect Basic version of the module manages the client connection balancing using a simple round robin balancing policy. SmartConnect Basic is limited to using static IP addresses, and to one IP address pool per external network subnet. The basic version is included with Isilon IQ's OneFS operating system as a standard feature.
- The SmartConnect Advanced version of the module offers CPU utilization, connection counting, and aggregate throughput client connection policies in addition to the simple round robin policy. The advanced version also allows IP address pools (known as *zones* in versions of OneFS earlier than 5.0) to be defined to support multiple DNS zones, and supports NFS failover.

To a client system, the cluster appears as a single network element. Both cluster and client performance can be enhanced when connections are more evenly distributed. SmartConnect provides intelligent connection balancing that does not require extensive configuration by users. Even in its minimum implementation, it can remove nodes that have gone offline from the request queue, and prevent new clients from mounting a down node. In addition, SmartConnect can be configured so new nodes are automatically added to the connection balancing pool.

Following are descriptions of the available SmartConnect DNS client connection balancing policies:

- **Round Robin:** This connection method works on a rotating basis, so that, as one node IP address is handed out, it moves to the back of the list; the next node IP address is handed out, and then it moves to the end of the list; and so on, depending on the number of nodes being used. This ensures that an orderly sequence occurs. This is the default state (once SmartConnect is activated) if no other policy is selected.



Note: Round Robin is the only connection policy included with SmartConnect Basic. The client connection balancing policies listed below are available only with SmartConnect Advanced.

- **CPU Utilization:** This connection method examines CPU use on each node, and then attempts to distribute the connections to balance the workload evenly across all nodes in the cluster.
- **Connection Count:** In this algorithm, the number of established TCP connections is determined, and then an attempt is made to balance connections evenly per node.
- **Network Throughput:** This method relies on an evaluation of the overall throughput volume, and then client connection balancing policies are set based on balancing this volume.

SmartConnect Basic configuration

SmartConnect Basic manages client connection balancing using a simple round robin balancing policy.

SmartConnect Basic is limited to using static IP addresses, and to one IP address pool per external network subnet. The basic version is included with Isilon IQ's OneFS operating system as a standard feature.

Simple round robin client connection balancing works on a rotating basis, so that, as one node IP address is handed out, it moves to the back of the list; the next IP address is handed out, and then it moves to the end of the list; and so on, depending on the number of nodes being used. This ensures that an orderly sequence occurs.



Note: SmartConnect requires that a new name server (NS) record be added to the existing authoritative DNS zone that contains the cluster. See [Add a name server record to the DNS infrastructure](#) on page 17 for details.

Configure connection balancing settings

You can configure connection balancing for your cluster's external network connections with SmartConnect.

Prerequisite: Before you can configure SmartConnect, you must first enable it by setting up a SmartConnect service address on the external network subnet that will answer incoming DNS requests.

You may have already configured SmartConnect while setting up an external network subnet using the Subnet Wizard. However, you can configure or modify SmartConnect's connection balancing settings at any time as your networking needs evolve.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **External Network Settings**, click the link for the subnet that you want to configure to use connection balancing.
The **Edit Subnet** page appears.
3. Under **Settings**, verify that the **SmartConnect service IP** has been configured; this is required to enable SmartConnect.

If that options says **Not set**, you must click **Edit**, and then select the specify the IP address that DNS requests will be directed to.

4. In the **SmartConnect settings** area , click **Edit**.
The **Configure Pool SmartConnect Settings** dialog box appears.
5. In the **Zone name** box, type a name for the SmartConnect zone that this IP address pool represents. The zone name must be unique among the pools served by the SmartConnect service subnet specified in Step 7 below.
6. In the **Connection policy** list, select the type of connection balancing policy the IP address pool for this zone will use. The policy determines how SmartConnect distributes incoming DNS requests across the members of an IP address pool.

- **Round Robin:** This method selects the next available node on a rotating basis. This is the default state (once SmartConnect is activated) if no other policy is selected.



Note: Round robin is the only connection policy available with SmartConnect Basic. The following client connection balancing policies are only available with SmartConnect Advanced.

- **Connection Count:** This method determines the number of open TCP connections on each available node to select which node the client will connect to.
 - **Network Throughput:** This method uses the overall average throughput volume on each available node to optimize the cluster usage.
 - **CPU Usage:** This method examines average CPU utilization on each node to optimize the cluster usage.
7. In the **SmartConnect service subnet** list, select the name of the external network subnet whose SmartConnect service will answer DNS requests on behalf of the IP address pool. A pool can have only one SmartConnect service answering DNS requests. If this option is left blank, the IP address pool it belongs to will be excluded when SmartConnect answers incoming DNS requests for the cluster.

If you have purchased the optional SmartConnect Advanced module, complete the following steps for the options under **SmartConnect Advanced**.

8. In the **IP allocation method** list, select the method by which IP addresses are assigned to the member interfaces for this IP address pool.
 - **Static:** Select this IP allocation method to assign IP addresses when member interfaces are added to the IP pool. As members are added to the pool, this method allocates the next unused IP address from the pool to each new member. Once an IP address is allocated, the pool member keeps the address indefinitely unless:
 - the member interface is removed from the network pool
 - the member node is removed from the cluster
 - the member interface is moved to another IP address pool
 - **Dynamic:** Select this IP allocation method to ensure that all IP addresses in the IP address pool are assigned to member interfaces, which allows clients to connect to any IP addresses in the pool and be guaranteed a response. If a node or an interface becomes unavailable, its IP addresses are automatically moved to other available member interfaces in the pool.

If you select the dynamic IP allocation method, you can further customize the dynamic behavior of the IP address pool in the following steps by specifying the SmartConnect **Rebalance policy** and the **IP failover policy** in the next two steps.

9. Select which type of SmartConnect **Rebalance policy** to use when IP addresses are redistributed when node interface members in an IP address pool become available after a period of unavailability:
 - **Automatic Failback:** This policy (the default) automatically redistributes IP addresses. The automatic rebalance is triggered by a change to:
 - the cluster membership
 - the cluster's external network configuration
 - a member network interface

- **Manual Failback:** This policy does not redistribute IP addresses until you manually issue a rebalance command via the command line or the web administration interface.
10. The **IP failover policy**—also known as NFS failover—determines how to redistribute the IP addresses among remaining members of an IP address pool when one or more members are unavailable. In order to enable NFS failover, you must first set the **IP allocation method** to **Dynamic**, and then select which type of **IP failover policy** to use:
- **Round Robin:** This method reassigns the IP address of a failed node interface member on a rotating basis.
 - **Connection Count:** This method determines the number of open TCP connections on each available node, and then attempts to proportionally redistribute the IP addresses from failed node interface members to the remaining members.
 - **Network Throughput:** This method evaluates the overall throughput volume, and then redistributes the IP addresses from failed node interface members to the remaining members based on optimizing this volume.
 - **CPU Usage:** This method examines average CPU utilization on each node, and then attempts to proportionally distribute the IP addresses from a failed node interface member to the remaining members across all nodes in the cluster.
11. Click **Submit** to save the SmartConnect settings.
The SmartConnect configuration changes are displayed on the **Edit Subnet** page.

Add a SmartConnect zone

You can add a SmartConnect zone to an external network subnet during the process of creating the subnet.

SmartConnect zones can be added during the creation of an external network subnet using the Subnet Wizard as described in [External network configuration](#) on page 18. These zones are synonymous with IP address pools.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **External Network Settings**, click **Add subnet** to start the **Configure network subnet** wizard.
The **Step 1 of 4 — Subnet Settings** wizard dialog box appears.
3. Configure the basic and advanced settings in the **Step 1 of 4 — Subnet Settings** wizard dialog box as described in [Add an external network subnet](#) on page 18, and then click **Next**.
The **Step 2 of 4 — IP Address Pool Settings** wizard dialog box appears.
4. Configure the IP address pool settings in the **Step 2 of 4 — IP Address Pool Settings** wizard dialog box as described in [Configure an IP address pool](#) on page 19, and then click **Next**.
The **Step 3 of 4 — SmartConnect Settings** wizard dialog box appears.
5. Configure the SmartConnect zone as described in [Configure SmartConnect settings](#) on page 20, and then click **Next**.
The **Step 4 of 4 — IP Address Pool Members** wizard dialog box appears.
6. Configure the IP address pool member settings as described in [Select the interface members for an IP address pool](#) on page 21, and then click **Submit**.
Your new SmartConnect zone is now configured as part of the cluster's new external network subnet.

Modify a SmartConnect zone

As your cluster's networking needs evolve, you may need to modify the settings of a SmartConnect zone that you created for an external network subnet using the Subnet Wizard.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Click the name of the external network subnet that contains the SmartConnect zone you want to modify.
The **Edit Subnet** page for the selected subnet appears.
3. In the **SmartConnect settings** area, click **Edit**.
The **Configure Pool SmartConnect Settings** dialog box appears.
4. Modify any basic or advanced SmartConnect settings, and then click **Submit**.
The SmartConnect configuration changes are displayed on the **Edit Subnet** page.

Disable a SmartConnect zone

You can remove a SmartConnect zone from an external network subnet.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Click the name of the external network subnet that contains the SmartConnect zone you want to disable.
The **Edit Subnet** page appears.
3. In the **SmartConnect settings** area, click **Edit**.
The **Configure Pool SmartConnect Settings** dialog box appears.
4. To disable the SmartConnect zone, delete the name of the SmartConnect zone from the **Zone name** box and leave it blank.
5. Click **Submit** to disable the SmartConnect zone.
The SmartConnect configuration changes are displayed on the **SmartConnect settings** section of the **Edit Subnet** page.

SmartConnect Advanced configuration

SmartConnect Advanced is an optional software module that provides multiple client connection balancing policies for multiple zones.

Storage administrators can select the most appropriate connection balancing policy for their network environment:

- **Round Robin:** Regular round robin works on a rotating basis, so that, as one node IP address is handed out, it moves to the back of the list; the next node IP address is handed out, and then it moves to the end of the list; and so on, depending on the number of nodes being used. This ensure that an orderly sequence occurs. This is the default state (once SmartConnect is activated) if no other policy is selected.
- **CPU Usage:** This method examines CPU use in each node, and then attempts to distribute the connections to balance the workload evenly across all nodes in the cluster.
- **Connection Count:** In this algorithm, the number of established TCP connections is determined, and then an attempt is made to balance these connections evenly per node.
- **Network Throughput:** This method relies on an evaluation of the overall throughput volume, and then client connection balancing policies are set based on optimizing this volume.



Note: SmartConnect requires that a new name server (NS) record be added to the existing authoritative DNS zone that contains the cluster. See [Add a name server record to the DNS infrastructure](#) on page 17 for details.

Configure NFS failover for a SmartConnect zone

NFS failover provides high availability by redistributing IP addresses among node interfaces in an IP address pool when one or more interfaces are unavailable.

Prerequisite: You must have the optional SmartConnect Advanced connection balancing module licensed in order to enable NFS failover on your Isilon cluster. SmartConnect Basic, which is provided with Isilon IQ's OneFS operating system as a standard feature, does not support NFS failover.

With NFS failover, SmartConnect Advanced ensures that all of the IP addresses in the pool are assigned to an available node. When a node goes down, the dynamic access IP addresses of the node will be redistributed among the remaining available nodes. Subsequent NFS client connections to the dynamically assigned IPs will be directed to the node newly assigned the address.

NFS failover is simply the combination of enabling dynamic IP allocation and IP failover. If your cluster has SmartConnect Advanced installed, you may have already enabled NFS failover during the process of running the Subnet Wizard to configure your external networking settings. You can also modify your subnet settings at any time to enable NFS failover for selected IP address pools.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Click the name of the external network subnet that contains the SmartConnect zone for which you want to enable NFS failover.

The **Edit Subnet** page for the selected subnet appears.

3. In the **SmartConnect settings** area, click **Edit**.

If the subnet contains more than one IP address pool, you'll first need to click the plus (+) sign to expand the view of the pool settings.

The **Configure Pool SmartConnect Settings** dialog box appears.

4. Under **SmartConnect Advanced**, set the **IP allocation method** to **Dynamic**. NFS failover will not work if this option is set to **Static**.
5. In the **IP failover policy** list, select which type of NFS failover method you want to use for the IP pool:
 - **Round Robin:** This method (the default) reassigns the IP address of a failed node interface member on a rotating basis.
 - **Connection Count:** This method determines the number of open TCP connections on each available node, and then attempts to evenly redistribute the IP addresses from failed node interface members to the remaining members.
 - **Network Throughput:** This method evaluates the overall throughput volume, and then redistributes the IP addresses from failed node interface members to the remaining members based on optimizing this volume.
 - **CPU Usage:** This method examines average CPU utilization on each node, and then attempts to evenly distribute the IP addresses from a failed node interface member to the remaining members evenly across all nodes in the cluster.
6. Click **Submit**.

The SmartConnect configuration change you made appears on the **Edit Subnet** page.

Configure dynamic IP allocation for a SmartConnect zone

Dynamic IP allocation is a feature of SmartConnect Advanced that ensures that all IP addresses are assigned to the nodes assigned to the IP address pool.

IP address allocation controls how OneFS assigns IP addresses to the node interfaces that belong to an IP address pool. This allows clients to connect to any IP addresses in the pool and be guaranteed a response. If a node or an interface becomes unavailable, its IP addresses are automatically moved to other available node interfaces in the pool.

Dynamic IP allocation has the following advantages:

- It enables NFS failover, which provides continuous NFS service on a cluster.
- It provides high availability because dynamic IP addresses are available to clients at all times.

If your cluster has SmartConnect Advanced installed, you may have already enabled dynamic IP allocation during the process of running the Subnet Wizard to configure your external networking settings. You can also modify your subnet settings at any time to enable dynamic IP allocation for selected IP address pools.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Click the name of the external network subnet that contains the SmartConnect zone you want to modify.
The **Edit Subnet** page for the selected subnet appears.
3. In the **SmartConnect settings** area, click **Edit**.
The **Configure Pool SmartConnect Settings** dialog box appears.
4. In the **IP allocation method** list, select **Dynamic** as the method by which IP addresses are assigned to the member interfaces for this IP address pool, and then click **Submit**.
The SmartConnect configuration change you made is displayed on the **Edit Subnet** page.

Configure a connection rebalancing policy for a SmartConnect zone

IP rebalancing is a feature of SmartConnect Advanced that controls how IP addresses are redistributed when node interface members for a given IP address pool become available again after a period of unavailability.

IP rebalancing is automatically enabled when the **IP allocation method** on the **Edit subnet** page is set to **Dynamic**.

If your cluster has SmartConnect Advanced installed, you may have already enabled an IP rebalancing policy during the process of running the Subnet Wizard to configure your external networking settings. You can also modify your subnet settings at any time to enable IP rebalancing for selected IP address pools.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Click the name of the external network subnet that contains the SmartConnect zone you want to modify.
The **Edit Subnet** page for the selected subnet appears.
3. In the **SmartConnect settings** area, click **Edit**.
The **Configure Pool SmartConnect Settings** dialog box appears.
4. Under **SmartConnect Advanced**, select which type of **Rebalance policy** to use when IP addresses are redistributed when node interface members in an IP address pool become available after a period of unavailability:
 - **Automatic Failback:** This policy (the default) automatically redistributes IP addresses. The automatic rebalance is triggered by a change to:
 - the cluster membership
 - the cluster's external network configuration
 - a member network interface
 - **Manual Failback:** This policy does not redistribute IP addresses until you manually issue a rebalance command via the command line or from the web administration interface.
5. Click **Submit**.
The SmartConnect configuration change you made appears on the **Edit Subnet** page.

NIC aggregation

Network interface card (NIC) aggregation, also known as link aggregation, is an optional IP address pool feature that allows you to combine the bandwidth of a node's physical network interface cards into a single logical connection that provides improved network throughput and redundancy.



Caution: Configuring link aggregation requires advanced knowledge of how to configure network switches. Consult your network switch documentation before configuring your cluster for link aggregation.

NIC aggregation can be configured during the process of creating a new external network subnet using the Subnet Wizard, or it can be configured on an existing subnet's IP address pool. When configuring a node via the web administration interface to enable NIC aggregation, storage administrators must be aware that:

- OneFS provides support for multiple link aggregation methods. Refer to [Configure NIC aggregation](#) on page 33 for details.
- Some NICs may only allow aggregation of ports on the same network card.
- For LACP and FEC aggregation modes, the switch must support IEEE 802.3ad link aggregation. Since the trunks on the network switch must also be set up, the Isilon node must be correctly connected with the right ports on the switch.
- A node's external interfaces cannot be used by an IP address pool in both an aggregated configuration and as individual interfaces. You must remove a node's individual interfaces from all pools before configuring an aggregated NIC, or else the web administration interface will display an error message when you try to save the configuration.



Note: You should enable link aggregation on the cluster before you enable it on the switch. If the cluster is configured but the switch is not configured, then the cluster can continue to communicate. If the switch is configured, but the cluster is not configured, it cannot communicate, and you are unable to configure the cluster for this purpose.

Configure NIC aggregation

You can configure cluster IP address pools to use NIC aggregation.

This procedure describes how to configure network interface card (NIC) aggregation for an IP address pool belonging to an existing subnet. You can also configure NIC aggregation while configuring an external network subnet using the Subnet Wizard.

Configuring NIC aggregation means that multiple, physical external network interfaces on a node are combined into a single logical interface. That means if a node has two external Gigabit Ethernet interfaces, both will be aggregated. Similarly, on a node with both Gigabit and 10 Gigabit Ethernet interfaces, its two 10 Gigabit Ethernet interfaces can be aggregated, as could its two Gigabit Ethernet interfaces. However, NIC aggregation cannot be used with mixed interface types.



Caution: Configuring link aggregation requires advanced knowledge of how to configure network switches. Consult your network switch documentation before configuring your cluster for link aggregation.



Note: You should enable link aggregation on the cluster before you enable it on the switch. If the cluster is configured but the switch is not configured, then the cluster can continue to communicate. If the switch is configured, but the cluster is not configured, it cannot communicate, and you are unable to configure the cluster for this purpose.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **External Network Settings**, click the name of the subnet that contains the IP address pool to which you want to add interface members.
The **Edit Subnet** page appears.
3. In the **Pool members** area, click **Edit**.



Note: In the case of multiple **IP Address Pools**, expand the pool to which you want to add the interfaces, then click **Edit** in the **Pool members** area.

The **Configure Pool Interface Members** dialog box appears.

4. In the **Available interfaces** table, click the aggregated interface for the node, which is indicated by **Aggregation** in the **Type** column.
For example, if you want to aggregate the network interface card for Node 2 of the cluster, click the interface named **ext-agg, Node 2** under **Available interfaces**, and then click the right-arrow button to move the aggregated interface to the **Interfaces in current pool** table.
5. Choose the appropriate aggregation mode, that corresponds to the network switch settings, from the **Aggregation mode** drop-down.

OneFS supports the following link aggregation modes:

- **Link Aggregation Control Protocol (LACP):** Support for the IEEE 802.3ad Link Aggregation Control Protocol (LACP). Isilon recommends this method for switches that support LACP. This is the default mode for new pools.
- **Legacy FEC mode:** Support for compatibility with aggregated configurations in earlier versions of OneFS.
- **Etherchannel (FEC):** Support for Cisco Fast EtherChannel. Newer implementation of the Legacy FEC mode.
- **Active / Passive Failover:** All data transmits via the master port, which is the first port in the aggregated link. The next active port in an aggregated link will take over if the master port is not available.
- **Round-Robin Tx:** Balances outbound traffic across all active ports in the aggregated link and accepts inbound traffic on any port.



Note: Consult your network switch documentation for supported link aggregation modes.

6. Click **Submit**.
The aggregated NIC configuration appears under **Pool members** on the **Edit Subnet** page.

Modify NIC aggregation

As your cluster's networking needs evolve, you may need to modify the settings for an IP address pool's aggregated NICs.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.

- Under **External Network Settings**, click the name of the subnet that contains the IP address pool with the NIC aggregation settings you want to change.
The **Edit Subnet** page appears.
- In the **Pool members** area, click **Edit**.



Note: In the case of multiple **IP Address Pools**, expand the pool to which you want to add the interfaces, then click **Edit** in the **Pool members** area.

The **Configure Pool Interface Members** dialog box appears.

- Modify the IP pool's aggregated NIC settings as described in [Configure NIC aggregation](#) on page 33, and then click **Submit**.



Note: A node's external interfaces cannot be used by an IP address pool in both an aggregated configuration and as individual interfaces. You must remove a node's individual interfaces from the **Interfaces in current pool** table before configuring an aggregated NIC; otherwise, the web administration interface will display an error message when you click **Submit**.



Note: Refer to your network switch documentation for supported link aggregation modes.

The modified NIC aggregation setting appears under **Pool members** on the **Edit Subnet** page.



Note: You should enable link aggregation on the cluster before you enable it on the switch. If the cluster is configured but the switch is not configured, then the cluster can continue to communicate. If the switch is configured, but the cluster is not configured, it cannot communicate, and you are unable to configure the cluster for this purpose.

Remove an aggregated NIC from an IP address pool

If you decide to remove an aggregated NIC configuration from an IP address pool because your network environment has changed, you cannot simply delete the aggregated setting but must instead replace it with single-NIC settings in order for the node to continue supporting network traffic.

- On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
- Under **External Network Settings**, click the name of the subnet that contains the IP address pool containing the aggregated NIC setting you want to remove.
The **Edit Subnet** page appears.
- In the **Pool members** area, click **Edit**.
The **Configure Pool Interface Members** dialog box appears.
- Select the name of the node's aggregated NIC that you want to remove in the **Interfaces in current pool** table, and then click the left-arrow button to move it into the **Available interfaces** table.
- Select one or more individual interfaces for the node in the **Available interfaces** table, and then click the right-arrow to move them into the **Interfaces in current pool** table.
- When you have completed modifying the node interface settings, click **Submit**.



Note: A node's external interfaces cannot be used by an IP address pool in both an aggregated configuration and as individual interfaces. You must remove a node's individual interfaces from the **Interfaces in current pool** table before configuring an aggregated NIC; otherwise, the web administration interface will display an error message when you click **Submit**.

The interface setting appears under **Pool members** on the **Edit Subnet** page.



Note: You should enable link aggregation on the cluster before you enable it on the switch. If the cluster is configured but the switch is not configured, then the cluster can continue to communicate. If the switch is configured, but the cluster is not configured, it cannot communicate, and you are unable to configure the cluster for this purpose.

NIC and LNI aggregation options

This table displays a list of network interface card (NIC) and logical network interface (LNI) mappings, including aggregation configuration options, for Isilon hardware platforms.

The following list provides guidelines for interpreting the aggregation options.

- Isilon nodes support multiple network card configurations.
- LNI numbering corresponds to the physical positioning of the NIC ports as found on the back of the node, numbered from left to right.
- Aggregated LNIs are listed in the order in which they are aggregated at the time they are created.
- NIC names correspond to the network interface name as shown in command-line interface tools such as `ifconfig` and `netstat`.

LNI	NIC	Aggregated LNI	Aggregated NIC	Aggregated NIC (Legacy FEC mode)
ext-1 ext-2	em0 em1	ext-agg = ext-1 + ext-2	lagg0	fec0
ext-1 ext-2 ext-3 ext-4	em2 em3 em0 em1	ext-agg = ext-1 + ext-2 ext-agg-2 = ext-3 + ext-4 ext-agg-3 = ext-3 + ext-4 + ext-1 + ext-2	lagg0 lagg1 lagg2	fec0 fec1 fec2
ext-1 ext-2 10gige-1 10gige-2	em0 em1 cxgb0 cxgb1	ext-agg = ext-1 + ext-2 10gige-agg-1 = 10gige-1 + 10gige-2	lagg0 lagg1	fec0 fec1

VLAN management

Virtual LAN (VLAN) tagging is an optional external network subnet setting that enables a cluster to participate in multiple virtual networks.

A VLAN is a group of hosts that communicate as though they are attached to the same local-area network regardless of their physical location. Enabling a cluster to participate in a VLAN provides the following advantages:

- Multiple cluster subnets can be supported without multiple network switches, so that one physical switch enables multiple virtual subnets.
- A VLAN provides increased security and privacy because network traffic across one VLAN is not visible to another VLAN.



Caution: Configuring a VLAN requires advanced knowledge of how to configure network switches to enable this option on a cluster. Consult your network switch documentation before configuring your cluster for a VLAN.

Configure VLAN tagging for a subnet

You can configure a cluster to participate in multiple virtual private networks, also known as virtual LANs or VLANs. You can also configure a VLAN during the process of creating a subnet using the Subnet Wizard.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.

2. Under **External Network Settings**, click the link for the subnet that contains the IP address pool to which you want to add interface members.
The **Edit Subnet** page appears.
3. In the **Settings** area for the subnet, click **Edit**.
The **Configure Subnet** dialog box appears.
4. In the **VLAN tagging** list, select **Enabled**.
5. In the **VLAN ID** box, type a number between **2** to **4094** that corresponds to the ID number set on the switch.
6. Click **Submit**.
The new VLAN settings are displayed on the **Edit Subnet** page.

Disable VLAN tagging for a subnet

If your network environment changes and you don't want a subnet to participate in a VLAN, you can disable that setting.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **External Network Settings**, click the link for the subnet that contains the IP address pool to which you want to add interface members.
The **Edit Subnet** page appears.
3. In the **Settings** area for the subnet, click **Edit**.
The **Configure Subnet** dialog box appears.
4. In the **VLAN tagging** list, select **Disabled**, and then click **Submit**.
The disabled VLAN setting appears on the **Edit Subnet** page.

Node provisioning rules

Provisioning rules automate the process of configuring the external network interfaces for new nodes when they are added to a cluster. This provides a convenient and consistent way to define in advance how the cluster treats new nodes as they're brought online.

When you initially configure a cluster, OneFS automatically creates a default provisioning rule named *rule0*. This default rule is configured so that all newly added nodes have their *ext-1* external network interface added to the default subnet named *subnet0* and use the default IP address pool named *pool0*.

You can modify this default provisioning rule to suit your networking needs. You can also add more provisioning rules to handle how specific types of nodes are configured when they are added to a cluster.

For example, you could create one provisioning rule governing how new Isilon storage nodes are configured, and another rule for new accelerator nodes. Similarly, you could also create provisioning rules specifically for Isilon's X-series storage and accelerator nodes.

Once the provisioning rules are configured, when a new node is added to the cluster it is automatically evaluated against the rules. If the new node's type (storage, accelerator, storage-x, or accelerator-x) matches that defined in a rule, then the new node's interface name is added to the subnet and the IP address pool specified in the rule. If there are multiple provisioning rules the new node is evaluated against each rule in turn until a match is found.



Note: Before creating a provisioning rule, you should verify that the IP address pool included in the rule has sufficient available IP addresses to accommodate the new node's client connections. If pool runs out of IP addresses, it will generate a log message.

Configure a provisioning rule

Configure one or more provisioning rules to automate the process of adding new nodes to your Isilon cluster.

Prerequisite: Before you can configure a provisioning rule, your cluster's external network settings—specifically the subnets and IP address pools—must first be configured. Before creating a provisioning rule, you should verify that the IP address pool included in the rule has sufficient available IP addresses to accommodate the new node's client connections.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **Provisioning Rules**, click **Add rule**.
The **Configure Network Provisioning Rule** dialog box appears.
3. In the **Name** box, type a unique name for the provisioning rule. The rule name can be up to 32 characters long and can include spaces or other punctuation.
4. In the **Description** box, optionally type a descriptive comment about the provisioning rule.
5. In the **If node type is** list, select which type of node you want the rule to apply to:
 - **All**: Applies the provisioning rule to all types of Isilon nodes that join the cluster.
 - **Storage**: Applies the provisioning rule only to Isilon IQ storage nodes that join the cluster.
 - **Accelerator**: Applies the provisioning rule to only Isilon IQ performance-accelerator nodes that join the cluster.
 - **Storage-X**: Applies the provisioning rule only to Isilon X-Series storage nodes that join the cluster.
 - **Accelerator-X**: Applies the provisioning rule only to Isilon X-Series performance-accelerator nodes that join the cluster.
6. In the **then assign interface** list, select one of the new node's interfaces to assign to the external network subnet and IP address pool specified in the rule:
 - **ext-1**: The primary external Gigabit Ethernet interface on the cluster.
 - **ext-2**: The secondary external Gigabit Ethernet interface on the cluster.
 - **ext-agg**: The primary and secondary external Gigabit Ethernet interfaces aggregated together.
 - **10gige-1**: The primary external 10 Gigabit Ethernet interface for Isilon X-series performance-accelerator nodes.
 - **10gige-2**: The secondary external 10 Gigabit Ethernet interface for Isilon X-series performance-accelerator nodes.
 - **10gige-agg-1**: The primary and secondary external 10 Gigabit Ethernet interfaces aggregated together.
7. In the **Subnet** list, select which external subnet the new node will join.
8. In the **Pool** list, select which IP address pool belonging to the subnet should be used by the new node.
9. Click **Submit**.
The new provisioning rule appears on the **Networking** page.

Modify a provisioning rule

After you have configured provisioning rules to automate the process of adding new nodes to your Isilon cluster, you can modify those rules as your networking needs evolve.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **Provisioning Rules**, click the name of the rule you want to modify.
The **Configure Network Provisioning Rule** dialog box appears for the selected rule.
3. Modify the provisioning rule settings as needed.
4. When you have finished modifying the provisioning rule, click **Submit**.
The modified rule appears on the **Networking** page.

Delete a provisioning rule

Provisioning rules automatically configure the external network interfaces for nodes that are added to a cluster. As your networking needs evolve, you may need to delete a provisioning rule that is no longer needed.

1. On the **Cluster** menu, click **Networking**.
The **Networking** page appears.
2. Under **Provisioning Rules**, click **Delete** next to the rule you want to delete.
A dialog box asks you to confirm that you want to delete the provisioning rule.
3. Click **Yes** to delete the rule, or click **No** if you want to retain it.

Cluster management

You can manage your cluster by viewing cluster and node status, generating email and SNMP alerts, and managing operations.

Cluster monitoring

You can view information about the health and performance of the entire cluster, including node status, client connections, new events, cluster size, cluster throughput, and CPU usage. You can view information for individual nodes, including node-specific network traffic, internal and external network interfaces, and drive status. You can configure real-time and historical performance to be graphed in the web administration interface.

You can use SNMP to remotely monitor hardware components (for example, fans, hardware sensors, power supplies, and disks) and software or subsystems for network-connected devices (for example, network interfaces, servers, switches, and CPU use). The Isilon cluster supports SNMP v1, v2c, and v3 for defined information. For monitoring assistance, Isilon provides Management Information Bases (MIBs) including:

- The base MIB module for Isilon Systems OneFS operating system
- Traps for Isilon products

The MIBs are located in `/usr/local/share/snmp/mibs`.

View cluster status

You can monitor the health and performance of a cluster in charts and tables that show the status and performance of individual nodes, client connections, events, cluster size, cluster throughput, and CPU usage.



Note: The **Cluster Status** tab displays **Monitoring** data in both **Current** and **Historical** views. The **Current** view reflects a recent moment of time or events as they are happening on the cluster, with frequent updating; the **Historical** view shows the state over a longer period of time, up the last two weeks of data. To view historical data, click **Current** in the drop-down menu in the upper-right corner of the **Monitoring** area, and then click **Historical**. The data in all the charts under **Monitoring** switch to historical view.

On the **Status** menu, click **Cluster Status**.

The **Cluster Status** tab appears and displays the following sections:

- **Status:** Displays health and performance statistics for each node of the cluster, including hard disk drive (HDD) and solid-state drive (SSD) usage. To view additional details about a particular node, click the node's ID number.
- **Client connection summary:** Displays a chart of the number of clients connected per node. To view a list of current connections, click the **Client Connections** tab.
- **New events:** Displays a list of notifications generated by system events, including the severity, unique instance ID, start time, alert message, and scope of the event. To view additional details about an event, click the **View details** link for the event.
- **Cluster size:** By default, the **Current** view displays a chart of used and available HDD and SSD space on the cluster, as well as any storage reserved for the virtual hot spare (VHS). The **Historical** view displays the total used space and cluster size over a one-year period.
- **Cluster throughput (file system):** By default, the **Current** view displays average inbound and outbound traffic volume passing through the nodes in the cluster over the past hour. The **Historical** view displays this information over the past two weeks. To view throughput statistics for a specific timeframe within the past two weeks, on the main menu, click **Status**, and then click **Throughput Distribution**.



Note: You can hide or show inbound or outbound throughput by clicking **Inbound** or **Outbound** in the chart legend. To view maximum throughput, next to **Show**, select **Maximum**.

- **CPU usage:** By default, the **Current** view displays average system, user, and total percentages of CPU usage over the past hour. The **Historical** view displays this information over the past two weeks.



Note: You can hide or show an individual plot by clicking **System**, **User**, or **Total** in the chart legend. To view maximum usage, next to **Show**, select **Maximum**.

View node status

You can view static and live information about node performance and activity. For each node in the cluster, you can view network interfaces, client connections, chassis and drive status, node capacity, node throughput, and CPU usage. You can also view detailed hardware status summaries.



Note: The **Node Status** tab displays **Monitoring** data in both **Current** and **Historical** views. The **Current** view reflects a recent moment of time or events as they are happening on the cluster, with periodic updating; the **Historical** view shows the state over a longer period of time, up to the last two weeks of data. To view historical data, click **Current** in the menu in the upper-right corner of the **Monitoring** area, and then click **Historical**. The data in all the charts under **Monitoring** switch to **Historical** view, with **Node capacity** appearing as a graph.

On the **Status** menu, click **Node Status**, and then click one of the node numbers. The **Node Status** tab appears.

You can view details of that node's status in these sections:

- **Node Status:** Shows details for the node, including model number, serial number, configuration, and uptime.
- **Network interfaces:** Shows information for each interface, including profile, address, status (whether it is active or has no carrier), and throughput.
- **Client connections:** Shows a list of clients connected to the node.
- **Chassis and drive status:** Shows the state of the drives in the node.

You can also click a bay or drive to view additional details. You can also perform drive operations on the **Drive Status** page.

- **Node capacity:** Shows the amount of data stored on the node.
- **Node throughput:** Shows the number of bytes written to and read from the OneFS file system in the past hour.
- **CPU usage:** Shows the CPU usage of the node, both graphically and numerically.



Note: Another way to display the **Node Status** tab is to click **Cluster** on the main menu, and then click **Cluster Status**. The **Cluster Status** tab appears. To view information for a particular node on the **Node Status** tab, click the node's number.

Monitor cluster size

You can view the amount of data stored on the cluster in both **Current** and **Historical** modes.

1. On the **Status** menu, click **Cluster Status**. The **Cluster Status** page appears.
2. On the right side of the page, under **Monitoring**, review the **Cluster size** information.

The chart provides a visual representation of the percentages of used and free space on the cluster. If you move the pointer over the chart, you can view the percentages of used and free space. By default, the chart shows the amount of used and free space for the labels.



Note: To view historical data for cluster size, click **Current** in the upper-right corner of the panel, and then click **Historical**. The data in the **Monitoring** section switches to historical view.

Monitor cluster throughput

You can view cluster throughput, graphically and numerically, for average and maximum usage. You can view information about the input and output traffic to and from the cluster's file system, as well as throughput distribution across the cluster.

1. On the **Status** menu, click **Cluster Status**.
The **Cluster Status** page appears.
2. On the right side of the page, under **Monitoring**, review the information in the **Cluster throughput** area.
The graph displays a summary of the inbound and outbound traffic, measured in bits per second, on the cluster. The values displayed refresh continually at regular intervals.
3. To view peak usage instead of **Average** usage, select **Maximum** at the bottom of the **Monitoring** area.
The graphs refresh to show the maximum usage numbers.



Note: To view historical data for cluster throughput, click **Current** in the upper-right corner of the Monitoring area, and then click **Historical**. The data in the **Monitoring** area switches to historical view.

To view throughput distribution using an alternative method

1. On the **Status** menu, click **Throughput Distribution**.
The **Throughput Distribution** tab appears, showing instantaneous results for all nodes through the default **Current** view.
2. For a historical view, in the top-right corner, click **Current**, and then click **Historical**.

A chart appears with a line representing the breakdown by node for the time period specified by the controls above the chart.

By default, the chart shows results from the last two weeks, the maximum range that can be displayed. To show data from a certain time period, specify the desired beginning date and time, specify how much time from that point to show in the chart, and then click **Show**. The chart updates to reflect your query. Depending on the size of the cluster and current cluster activity, it may take several minutes to update the chart.

Monitor CPU usage

You can view current and historical CPU usage by the system and the users on the cluster, graphically and numerically.

1. On the **Status** menu, click **Cluster Status**.
The **Cluster Status** page appears.
2. In the upper-right area of the page, under **Monitoring**, view the chart in the **CPU usage** panel. The chart displays percentages of CPU usage over time, by the system for its operations and users on the cluster.
 - To view peak usage instead of **Average** usage, click **Maximum** at the bottom of the **Monitoring** section. The chart refreshes to show peak CPU usage statistics.
 - To view historical data for CPU usage, click **Current** in the upper-right corner of the section, and then click **Historical**. The data in the charts under **Monitoring** switch to historical view.

View client connections

You can view client connections on the cluster, organized by the protocol used and the individual node. You can search for and sort connections by client.

1. On the **Status** menu, click **Cluster Status**.
The **Cluster Status** page appears. The chart displays smb, NFS, FTP, and HTTP client TCP connections to all nodes in the cluster.
2. In the left column, review the **Client connection summary** information, which is organized by individual node.

- To view a list of current connections organized by protocol and node, click the **Client Connections** tab.
- To view specific clients, type the **Client** name, select a **Protocol** from the list, select a particular **Node** (if applicable), and then click **Search**. (For the **Client** name, you can enter wildcard searches as well, for example "10.*.") The results appear in the results box.

View cluster and node statistics

The `isi statistics` command-line tool provides an interface for querying and displaying Isilon cluster performance and usage statistics.

Several modes of operation provide methods for viewing different types of statistics. You can use generic and type-specific options to control statistics filtering, aggregation, and reporting for each mode of statistics reporting. You can access these modes of operation by typing subcommands in the `isi statistics` tool.

For more information about the generic statistics options, descriptions of the individual subcommands, and options specific to these subcommands, search for "isi statistics command-line tool" in the Isilon Knowledge Center.

View active alerts

View notifications of active alerts generated by system events.



Note: Clicking **System alert** next to the icon of the yellow exclamation point in the upper-right corner of the main menu opens the **Historical Alerts** page, where you can view current alert activity and quiet alerts.

On the **Status** menu, click **Cluster Status**. The **Cluster Status** page appears.

- To view all alerts, click **View all alerts**. The **Alert History** page appears, organized by **Active Alerts** and **Historical Alerts**.



Note: Each alert entry displays the information included on the **Cluster Status** tab, including date and time, priority level, alert message, ID number, and quiet status.

- To display the **Alert History** page, on the main menu, click **Status**, and then click **Alerts**. The **Alert History** page appears, organized by **Active Alerts** and **Historical Alerts**.



Note: The bottom of the page contains an **Active alerts** section with columns indicating the date and time the alert was generated, the priority level of the alert, the alert message, the unique alert event ID for tracking, and quiet status.

View cluster logs

You can view detailed log data and recent log activity. These logs provide information about recent cluster status and performance parameters.

1. On the **Status** menu, click **Logs**.
The **Logs** page appears.
2. For the node number containing the logs you want to view, click the link for the specific log type. For example, click **messages** to view message logs.
The **View Log** page appears, displaying the details of the log file.

SNMP monitoring

You can use SNMP to remotely monitor hardware components (for example, fans, hardware sensors, power supplies, and disks) and software or subsystems for network-connected devices (for example, network interfaces, servers, switches,

and CPU use). You can enable SNMP monitoring on individual nodes on your cluster. You can also monitor cluster information from any node.



Note: This page includes input examples for running `snmpwalk` on a cluster. Your version of this program might differ or require different arguments.

Isilon recommends that you configure your network monitoring system (NMS) to query each node directly. This approach ensures that all nodes have external IP addresses and therefore can respond to SNMP queries.

Because the default setting for SNMP proxying is **enabled**, the SNMP implementation on each node is configured automatically to proxy for all other nodes in the cluster except itself. This proxy configuration allows the Isilon Management Information Base (MIB) and all available standard MIBs to be exposed seamlessly through the use of context strings for SNMP v3 or community strings for SNMP v2c.



Note: SNMP proxying is not available for SNMP v1.

Enabling SNMP proxying requires you to append “_node_” followed by the logical node number for the node to the community/context string. For example, in a three-node cluster, if node 1 has an IP address of 10.54.145.49, and the community string is set to “public,” you can query node 2 through that node by using a community string “public_node_2”, such as:

```
snmpwalk -v 2c -c public_node_2 10.54.145.49 sysLocation.0
```

The `net-snmp` command causes node 1 at 10.54.145.49 to proxy the request to node 2. For SNMP v3, use this modified community string as the context string. For example:

```
snmpwalk -v 3 -n public_node_2 -l AuthNoPriv -A password 10.54.145.49 sysLocation.0
```



Note: For Nagios users, Isilon provides two downloadable Nagios configuration files based on your cluster configuration. The *Global Nagios Configuration* document contains a generic configuration that works for any cluster. The *Cluster Nagios Configuration* document contains a Nagios configuration specific to this cluster.

Configure SNMP monitoring

You can configure your cluster to perform SNMP monitoring.

1. On the **Cluster** menu, point to **Cluster Settings**, and then click **SNMP Monitoring**. The **SNMP Monitoring** page appears.
2. In the **Service** area, disable or enable SNMP monitoring as needed. The default setting is **SNMP is currently enabled**.
 - To disable SNMP monitoring, click **Disable**. The SNMP process is disabled for the entire cluster.
 - To enable SNMP monitoring, select **Enable**. The SNMP process is enabled for the entire cluster.
3. In the **Settings** area, under **General Settings**, configure settings as needed.
 - a. To enable or disable the version or versions of the SNMP protocol to which the nodes respond, next to **Protocol access**, select a protocol choice from the list of three options.

The security model is different for each of the three SNMP protocol versions: v1, v2c, and v3. However, because the security-model implementations for v1 and v2c are the same, these protocol versions are controlled together. SNMP v3 access is enabled or disabled independently of v1 and v2c.

 - **Allow access via SNMP v1 and SNMP v2c only**
 - **Allow access via SNMP v3 only**
 - **Allow access via SNMP v1, SNMP v2c and SNMP v3**
 - b. In the **System location** box, type the system location name.

This setting is the value that the node reports when responding to queries for the OID `SNMPv2-Mib::sysLocation.0`. Type any name that helps to identify the location of the node.

- c. In the **System contact** box, type identifying information (for example, system-contact name or email address) for the person or entity that manages the cluster.

This setting is the value the node reports when responding to queries for the OID
`SNMPv2-MIB::sysContact.0`.

4. In the **Settings** area, configure protocol settings as needed.



Note: You must enable v1/v2c protocols in the **General Settings** section to ensure that the settings are active.

- In the **Settings** area, under **SNMP v1/v2c Settings**, type the name of the community to which you are granting read-only rights. The default setting is *public*.



Note: You must enable v1/v2c protocols in the **General Settings** section to ensure that the settings are active.



Note: OneFS does not support writable OIDs; therefore, no write-only community string setting is available.

- In the **Settings** area, under **SNMP v3 Settings**, choose a user with read-only permissions and choose an SNMP v3 password.

To change the name of the user with read-only permissions, in the **Read-only user** box, type the SNMP v3 security name. Do not include any space characters in the **Read-only user** setting. The default read-only user is "general."

If you want to set a new SNMP v3 authentication password, in the **SNMP v3 password** box, type the new password for the read-only user. The password must contain at least eight characters, and must not contain any space characters. The default password is "password."

In the **Confirm password** box, retype the SNMP v3 authentication password.



Note: You must enable v3 protocols in the **General Settings** section to ensure that the settings are active.

5. Click **Submit**.

Download MIBs

Management Information Base (MIB) documents define human-readable names for managed objects, and specify their data types and other properties. You can download MIBs for SNMP monitoring.

The Isilon MIBs serve two purposes: to augment the information available in standard MIBs, and to provide Isilon-specific information that is not available in standard MIBs.

1. On the **Cluster** menu, point to **Cluster Settings**, and then click **SNMP Monitoring**.
The **SNMP Monitoring** page appears.
2. In the **Downloads** area, click the appropriate download link.
The **Opening...** dialog box appears.
3. Save the selected MIB.
 - To save the selected MIB in Notepad or another application, click **Open with**. To open the file with Notepad, click **OK**. The file opens in Notepad.
 - To save the MIB to your hard drive, click **Save File**. To open the file with another application, select **Other**. The **Choose Helper Application** dialog box appears. Click one of the applications in the list and then click **OK**. The **Opening...** dialog box appears. If you want to use a different application than the list that appears, click **Browse** to navigate to the application, select the application, and then click **Open**.

The file opens in the application.

Cluster configuration

You can manage the cluster configuration by adding or removing a node, setting the join mode, upgrading the cluster operating system, shutting down or rebooting the cluster, and setting the date, time, and name of the cluster.

Configure the OneFS web Interface port

You can create a `webui_port.xml` file to configure the OneFS web interface to use a port other than port 8080.

1. Log in to the cluster.
2. Use `vi` to create a new file named `webui_port.xml` in the `/etc/mcp/override` directory.

```
# vi /etc/mcp/override/webui_port.xml
```

3. Insert the following xml code into the body of the file.

```
<?xml version="1.0" encoding="utf-8"?><isi-data file="webui-port">
  <modify-text id="port">8080</modify-text>
</isi-data>
```

4. Replace 8080 in the xml code with the new port number you want to use for the web interface.



Note: The web interface can use all non-negative integer port numbers between 1025 and 65535, except for port 8081 and port 8086.

5. Save the file, and close the editor.

The cluster will recognize the change and copy the file across all the nodes and reconfigure the web interface. The web interface will use the new port the next time you connect to the web interface.

Add nodes to a cluster

You can add nodes to your existing clusters through the web administration interface. From the **Add Node** page, you can export the configuration settings from an existing cluster to a new unconfigured node, and add the node to the cluster.



Note: “Upgrade” refers to the target OneFS version being equal to or later than the current version. “Downgrade” refers to the target version being earlier than the current version.

In OneFS 5.5 and later versions, the method you use to add new nodes (running any version of OneFS) to an existing cluster has been updated. The goal of the new upgrade functionality is two-fold:

1. To avoid the consequences of an update failure on adding a node (for example, the need to reset the node and smartfail it from the cluster).
2. To allow multiple unconfigured nodes to run upgrades simultaneously.



Important: The **Add Node** feature in OneFS 5.5 and later versions is not compatible with earlier implementations. To add nodes or clusters running versions earlier than 5.5 to nodes or clusters running 5.5 or later, you must manually reimage them to at least 5.5.

An additional benefit of this enhanced functionality is that you avoid having to unnecessarily upgrade a node being added to a patched cluster if the node already contains the patches.



Note: If you do not have available IP addresses (the number in the **Cluster Details** section is "0"), then you cannot add nodes to the cluster.

1. On the **Cluster** menu, point to **Cluster Management**, and then click **Add Node**. The **Add Node** page appears.
2. In the **Available Nodes** section of the page, locate the node that you want to add, and then click **Add node**. The **Status** and **Actions** columns update periodically to indicate the status of the Add Node operation.

The node to be added disappears from the list of available nodes once it has been successfully added, and the data in the **Cluster Details** section of the page is updated to show an increase in the number of nodes and in the cluster capacity and a decrease in the available IP addresses.

You can view any occurring errors that appear in the **Status** column, and you must confirm errors in the **Actions** column.

Remove nodes from a cluster

You can remove nodes from a cluster. When you remove a standard storage node from a cluster, the system automatically smartfails the node before removing it to ensure that the data the node contains is transferred to other nodes in the cluster. Non-storage Accelerator nodes are removed immediately when you initiate the removal.

Because removing a storage node from the cluster will delete all data from that node, the Isilon OneFS software automatically restripes the data to other nodes in the cluster using FlexProtect before actually removing the node. FlexProtect is responsible for protecting data in the cluster based on the configured protection policy. It detects and repairs files and directories that are in a degraded state.

1. On the **Cluster** menu, point to **Cluster Management**, and then click **Remove Node**.
The **Remove Node** page appears.
2. From the list of nodes on the page, click to select the node you want to remove from the cluster, and then click **Submit**.
The **Remove Node** page refreshes, and prompts you to confirm the removal.
3. Click **Confirm**.



Note: If you are removing a storage node, the system first smartfails the node to ensure that no data is lost. After a slight delay, the smartfail process begins and the **Cluster Status** page appears, displaying status information about the removed node. When the smartfail process is complete, the node no longer appears on the **Cluster Status** page.

Upgrade cluster operating system

OneFS provides the option to perform either a rolling or simultaneous upgrade of the cluster operating system.

A rolling upgrade individually upgrades and restarts each node in the cluster sequentially. A simultaneous upgrade installs the new operating system and restarts all nodes in the cluster at the same time. During a rolling upgrade, the cluster remains online and continues serving clients with no interruption in service. A Simultaneous upgrade requires a temporary interruption of service during the upgrade process.



Note: Rolling upgrades are available in OneFS 5.5.3 or later.

Before beginning a rolling upgrade, OneFS compares the current operating system to the new version to ensure that a rolling upgrade is permitted. If a rolling upgrade is not compatible with your system—typically due to incompatibility between certain OneFS versions—OneFS displays a warning message. If a rolling upgrade is not permitted, you can still perform a simultaneous upgrade of the cluster operating system.



Note: Rolling upgrades typically start with the node directly following the node from which you selected the upgrade. For example, if you begin a rolling upgrade from node 1, the update process begins with node 2. The upgrade process then proceeds sequentially through the other nodes in the cluster. The final node in the upgrade process is the node from which you started the upgrade process.

For specific instructions on how to upgrade the cluster operating system, see the *Isilon OneFS Release Notes*.

Shut down or reboot the cluster

You can shut down the cluster completely or reboot the cluster.

1. On the **Cluster** menu, point to **Cluster Management**, and then click **Shut Down Cluster**.
The **Shut Down Cluster** page appears.

2. To shut down the cluster, click **Shut down**, and then click **Submit**.
 - To shut down the cluster, click **Shut down**, and then click **Submit**.
The **Shut Down Cluster** page refreshes, and prompts you to confirm the shut-down selection. Click **Confirm**. A page appears and confirms the shut-down action.
 - To stop the cluster and then restart it, click **Reboot**, and then click **Submit**.
The system prompts you to confirm the reboot selection. Click **Confirm**.
The **Rebooting Cluster** page appears. When the reboot is complete, the web administration login page appears.

Set the cluster name

You can assign a friendly name to your cluster to make the cluster and its nodes more easily recognizable on your network. The name will be used in conjunction with a node number to identify each node in the cluster. For example, the first node in a cluster named “Images” might be named “*Images-1*.”

You must add the specified cluster name to your DNS servers.



Note: Valid cluster names must begin with a letter and can contain only numbers, letters, and hyphens.

1. On the **Cluster** menu, point to **Cluster Settings**, and then click **Cluster Identity**.
The **Cluster Identity** page appears.
2. In the **Cluster name** box, type the name that you want to assign to the cluster, and then click **Submit**.



Note: For Active Directory domains, limit the cluster name to 11 characters or less. Two to four characters are appended to the cluster name you specify to generate a unique name for each node.

Specify company contact information

To enable Isilon and recipients of your alert notifications to contact you, you can specify your company's contact information.

1. On the **Cluster** menu, point to **Cluster Settings**, and then click **Cluster Identity**.
The **Cluster Identity** page appears.
2. In the **Contact Information** section, specify the company name, location, primary and secondary contact names, email addresses, and primary and secondary phone numbers.
3. Click **Save**.

Specify email server settings

To send alert notifications as email messages, you must specify an SMTP mail sever. If your SMTP server is configured to use authentication, you can also enable SMTP authentication.

1. On the **Cluster** menu, point to **Cluster Settings**, and then click **Email Settings**.
The **Email Settings** page appears.
2. In the **SMTP relay address** box, type the domain name or IP address of your SMTP mail server, for example, *mail.company.com*.
3. In the **SMTP relay port** box, type the number of the TCP/IP port on which your SMTP mail server listens.
4. In the **Originator email address** box, type the primary contact's email address.
5. In the **Use SMTP AUTH** area, specify whether you want to enable or disable SMTP authentication:
 - To disable SMTP authentication, click **No**.
 - To enable SMTP authentication, click **Yes**.

6. If, in the previous step, you enabled SMTP authentication, configure additional SMTP authentication settings:
 - a. In the **Username** box, type the username for the SMTP server
 - b. In the **Password** box, type the password for the SMTP server.
 - c. Specify a **Connection security** setting:
 - To disable connection security, click **No security**.
 - To enable connection security, click **STARTTLS**.
7. Click **Submit**.

Manually set the cluster date and time

As an alternative to the NTP (Network Time Protocol) method, in which the cluster automatically synchronizes its date and time settings through an NTP server, you can manually set the date and time reported by the cluster.

1. On the **Cluster** menu, point to **Cluster Settings**, and then click **Date and Time**.
The **Date and Time** page appears, and displays a list of each node's IP address and the current date and time settings for each node.
2. Specify the cluster's date, time, and time-zone settings:
 - a. In the **Date and time** area, specify the month, date, year, hour, and minute settings.
 - b. In the **Time zone** list, click a time zone.
 - c. If your desired time zone is not on the list, in the **Time zone** list, choose **Advanced**.
The **Advanced time zone** list is enabled.
 - d. In the **Advanced time zone** list, click a more specific time zone.
3. Click **Submit**.

The **Set Date and Time** page refreshes, with a notification reporting the current date and time.



Note: Alternatively, you can configure the NTP service to ensure that all nodes in a cluster are synchronized to the same time source.

Set the cluster join mode

You can specify the methods by which nodes can be added to the cluster: The attach or join methods can be used in **Manual** mode, and the attach method can be used only in **Secure** mode.

With the attach method, the cluster invites the node to join the cluster. With the join method, the node makes a request to join the cluster.

1. On the **Cluster** menu, point to **Cluster Settings**, and then click **Join Mode**.
The **Join Mode** page appears.
2. Select the join mode in which you want the cluster to operate:
 - **Manual Mode:** If you set the cluster to use **Manual** mode, *either* join or attach can be used to add nodes to a cluster. You can use join mode when you have multiple Isilon clusters on the same network, but this method requires the installer to know which existing cluster the new node is to be joined to.
 - **Secure Mode:** If you enable **Secure** mode, you must use the attach method to add nodes rather than the join method. **Secure** mode is recommended when you have multiple Isilon clusters on the same network because it will prevent new nodes from being added to the wrong cluster. **Secure** mode prevents any new nodes from automatically joining the cluster.
3. Click **Submit**.

Cluster services

You can configure cluster services that manage telnet access, cluster encoding, and Network Time Protocol (NTP) settings.

Configure telnet

You can communicate with the cluster using the telnet protocol.

1. On the **Cluster** menu, point to **Cluster Settings**, and then click **Telnet (Remote Access)**.
The **Telnet (Remote Access)** page appears.
2. To turn the telnet service on or off:
 - Click **Disable** to turn off the telnet service. Telnet is disabled by default.
 - Click **Enable** to turn on the telnet service.
3. Click **Submit**.

Configure NTP

You can specify one or more Network Time Protocol (NTP) servers that synchronize the system time on the cluster. The cluster periodically contacts all servers and sets the date and time based on the information it receives.

For information on setting the time manually, see [Manually set the cluster date and time](#) on page 48.

1. On the **Cluster** menu, point to **Cluster Settings**, and then click **NTP**.
The **NTP** page appears.
2. In the **Server IP or hostname** box, type the host name or IP address of the NTP server, and then click **Add**.
The server appears in the list of NTP servers.



Note: To remove an NTP server, select the check box next to the server name in the **Server** area, and then click **Delete**.

3. Confirm the settings by clicking **Submit**.

Configure character set encoding

You can modify the character set encoding for the cluster after installation.



Note: UTF-8 is the default character set for Isilon IQ nodes, and will be selected unless changed by users. Also note that the cluster must be rebooted before changes in the character set encoding are applied.



Important: Character set encoding is usually established at installation of the cluster. Modifying character encoding settings at a later date can render files unreadable if done incorrectly. Modify settings only if necessary, after consultation with Isilon Systems support staff.

1. On the **File System** menu, point to **File System Settings**, and then click **Character Encoding**.
The **Character Encoding** page appears. The page indicates the default character set and the set that is currently being used.
2. To change the encoding used by the cluster, click the **Character encoding** drop-down box, and select from the list by highlighting your choice. The listed items are the only character sets supported.
3. Click **Save**.

The page will reload and present a confirmation message that the encoding will be changed after cluster reboot.



Note: The cluster must be rebooted for encoding changes to take effect. The cluster can be rebooted via the **Shut Down Cluster** page.

- On the **Shut Down Cluster** page, select **Shut down**, and then click **Submit**.
After the cluster has come back up, the encoding will be changed to your new selection.

Enable access time tracking

By default, the Isilon cluster does not track the times at which files are accessed. However, you can enable access time tracking if necessary.

As an example, you must enable access time tracking if you want to configure SyncIQ policy criteria that match files based on when they were last accessed.




Note: Enabling the access time tracking option can affect cluster performance. It is recommended that you enable this option only if necessary.

- On the **File System** menu, point to **File System Settings**, and then click **Access Time Tracking**.
The **Access Time Tracking** page appears.
- Select the **Enabled** option.
- In the **Precision** text box, specify (in **Seconds, Minutes, Hours, Days, Weeks, Months, or Years**) how often to update last-accessed times.
For example, if you were to configure a **Precision** setting of **1 Day**, the cluster would update the last-accessed time once per day, even if some files were accessed more often than once per day.
- Click **Submit**.

SAS drive LED status

SAS hard drives use an external LED to indicate the current status of the drive.

LED Behavior	Description
OFF	<p>The drive is in a standby or stopped state. The drive can be removed without damage to the drive.</p> <p> Note: A drive that has not physically failed—but which you will remove from the node for some other reason—must be smartfailed before removing the drive.</p>
ON	The drive is active and ready to receive a command.
Blinking	The drive is active and currently processing a command.

File system management

You can share files and directories over a network, configure authentication settings and user permissions to access those files, configure the file system data protection level, configure and view backup processes, and tune write performance. File System Explorer enables you to browse the file system to review and configure these settings on a per-file and per-directory basis.

File sharing

The Isilon clustered storage system supports file sharing via SMB for Windows, NFS for UNIX, secure shell (SSH), FTP, and HTTP. Because multi-protocol support is built into the OneFS operating system, a single directory can be configured to serve as an SMB share, NFS export, and document root directory.

Users who have proper credentials and privileges can create, modify, and read data using one of the supported file sharing protocols:

- **Windows file sharing (SMB).** Allows Microsoft Windows clients to access files stored on the Isilon cluster.
- **UNIX file sharing (NFS).** Allows UNIX, Linux, Mac OS X, Solaris, and other UNIX-based clients to access files stored on the Isilon cluster.
- **HTTP (with optional DAV).** Allows clients to access files from a web browser, such as Internet Explorer, Mozilla Firefox, Opera, or Safari.
- **FTP.** Allows any client equipped with an FTP client program (such as a command-line FTP client, FileZilla, Firefox, or Internet Explorer) to access files through the FTP protocol.

By default, only the SMB and NFS protocols are enabled on the cluster.

The root directory for all file system data in the cluster is `/ifs` (Isilon file system). The SMB protocol includes an `/ifs` share, and the NFS protocol includes an `/ifs` export.

The system also enables you to set Windows- and UNIX-based permissions.

For users working on Windows operating systems, the Isilon cluster supports the following authentication methods:

- **Anonymous mode.** Allows users to access files without providing any credentials.
- **User mode.** Allows file access to users from any configured authentication source.

For users working on UNIX-based operating systems, the Isilon cluster supports the following additional authentication methods:

- **NIS.** Network Information Service (NIS) is a client/server directory service protocol for distributing system configuration data such as user and host names between computers on a network. NIS allows users to authenticate with credentials that are stored in the directory services. You can use other implementations of NIS, as long as they are compatible with the Sun Microsystems implementation.
- **LDAP.** Lightweight Directory Access Protocol (LDAP) is an application protocol for querying and modifying directory services running over TCP/IP. LDAP is available from a variety of vendors, and allows users to authenticate with credentials that are stored in an LDAP repository.

For additional information about each of the supported authentication methods, see "Authentication."

Multi-protocol file sharing

OneFS supports SMB for Windows file sharing and NFS for UNIX file sharing. You can configure your cluster to use either protocol exclusively, or use them together for mixed Windows/UNIX environments. For each of these scenarios, HTTP, FTP, and SSH can also be enabled.

SMB (Windows file sharing protocol)

OneFS supports Server Message Block versions 1 and 2 (SMB1, SMB2) for mapping Windows network drives to the Isilon cluster. You can configure the rules and other settings that govern the interaction between your Windows network and individual SMB shares on the cluster.

A Windows file share named `/ifs` is configured and enabled by default on all Isilon clusters. Depending on the authentication mode being used, you can create share options as you would in a standard Windows environment.

The Isilon cluster supports the following modes for authentication through Windows:

- **Anonymous mode.** Allows users to access files without providing any credentials.
- **User mode.** Allows file access to users from any configured authentication source.

Windows shares provide access to file system resources over the network. After you add a share in user mode, you can grant permissions to users and groups to carry out operations such as reading, writing, and setting access permissions. Although you can directly configure advanced Windows network settings, it is recommended that you do so only if you have experience working with the Windows environment.

SMB protocol behavior

The SMB protocol uses security identifiers (SIDs) exclusively for authorization data. All identities are converted into SIDs during retrieval and back to their on-disk representation before storage.

A newly created file inherits the access control list (ACL) of its parent directory. If no inheritable ACL exists, a default ACL is created. This behavior can be controlled by a per-share configuration setting.

View SMB summary information

You can view the status of the SMB file sharing service and manage any shares that have been created on the cluster.

On the **File Sharing** menu, point to **SMB**, and then click **Summary**.

The **SMB > Summary** page appears, and displays the following sections:

- **SMB Service:** Displays the status of the SMB file sharing service and the selected security mode.
- **Shares:** Displays existing SMB shares and provides the ability to modify, delete, and add new shares to the cluster.

Configure SMB settings

You can enable or disable the SMB service, set the preferred security mode, and change security and default share settings.

1. On the **File Sharing** menu, point to **SMB**, and then click **Settings**.

The **SMB Settings** page appears.

2. View or change the settings in the following sections:

- **File Sharing Service:** Displays options for the SMB file sharing service and security mode.

Setting	Description
Service status	Enables or disables the SMB service.
Security mode	Sets the security mode to Anonymous or User.



Caution: You should not configure the remaining settings below unless you have experience working with SMB. Changes to these parameters may impact the availability of the SMB file sharing service.

- **General Settings:** Displays general options for the server.

Setting	Description
Server string	A descriptive name for SMB shares on the server.

- **Security Settings:** Displays the security settings for SMB.

Setting	Description
Enable security signatures	Indicates if packet-signing is allowed.
Guest user	Specifies the fully-qualified user name that will be used for guest access.
Require security signatures	Indicates if all packets must be signed.

- **Default Share Settings:** Displays the default settings to be used when a share is created.



Note: These parameters apply to all shares but will be overridden by settings that are specified at the share level.

File and Directory Permissions

Setting	Description
Directory create mask	Defines which mask bits are applied when a directory is created.
Directory create mode	Defines which mode bits are applied when a directory is created.
File create mask	Defines which mask bits are applied when a file is created.
File create mode	Defines which mode bits are applied when a file is created.
Create permissions	Indicates the default source permissions to apply when a file or directory is created.

Performance Parameters

Setting	Description
Change notify	Configures notification of clients when files or directories change, which helps prevent clients from seeing stale content but requires server resources.
Oplocks	Indicates if an oplock request is allowed. An oplock, or opportunistic lock, allows clients to provide performance improvements by using locally cached information.
Strict locking	Indicates if the server checks for and enforces byte-range locks on read and write operations.

Security Parameters

Setting	Description
Allow delete read only	Indicates if read-only files can be deleted.

Setting	Description
Allow execute always	Indicates if a user with read access to a file can also execute the file.
Hide dot files	Indicates if UNIX configuration files, or dot files, are hidden.
Impersonate guest	Determines guest access to a share.
Impersonate user	Allows all file access to be performed by a specific user. this must be a fully qualified user name.

- Click **Submit** or **Cancel**.



Note: Changes are not committed until you click **Submit**. Clicking **Cancel** will exit the SMB **Settings** page immediately without committing any changes.

SMB share management

SMB shares provide Windows client access to files and directories on the cluster. OneFS supports %U and %D variable expansion and automatic provisioning of user home directories.

Add an SMB share

When adding or modifying an SMB share, you can override default permissions, performance, and access settings. You can configure SMB home directory provisioning by using directory path variables to automatically create and redirect users to their own home directories.



Note: To configure SMB home directory provisioning, the following configuration settings are required:

- The **Directory to share** path must contain the %U variable for user name substitution, and optionally %D for the domain name (to prevent conflicts for users with the same user name). These variables are case-sensitive.
- Both the **Allow Username Expansion** and **Automatically Create User Directory** options must be selected.

- On the **File Sharing** menu, point to **SMB**, and then click **Add Share**.

The SMB **Add Share** page appears.

- In the **Name and Path** area, in the **Share name** box, type a name for the share.

The share name can include up to 80 characters. The following characters are not supported: \ / ? * [] : ; | = , + < > " { }

- In the **Description** box, type a descriptive comment about the share.

A description is optional, but can be helpful when managing multiple shares.

- In the **Directory to share** box, type the full path of the share, beginning with /ifs, or click **Browse** to locate the share.



Note: You can use the %U and %D variables as substitutions for the user's Windows user name and domain name / workgroup, respectively. For example, for a user in a domain named DOMAIN whose username is user_1, the path /ifs/home/%D/%U is interpreted as /ifs/home/DOMAIN/user_1.

If the directory path contains one or both of these variables, the following additional settings are available:

- Allow Username Expansion:** Expands variables in the directory path. %D expands to the user's Windows domain name, normalized to uppercase. %U expands to the user's Windows name, normalized to lowercase. If this option is not selected, %D and %U are treated as regular text rather than variables.
- Automatically Create User Directory:** When %U and %D are expanded, if the resulting directory path does not already exist, this setting creates a home directory for the user. To use this setting, you must also select the **Allow Username Expansion** option. If username expansion is enabled but automatic directory creation is not, any attempt to connect to an expanded path that does not already exist will simply fail.

5. Specify a Directory ACLs setting:

- To apply a default ACL (similar to the permissions in a Windows file server) to the shared directory, click the **Apply Windows default ACLs** option. This is a one-time operation, applicable only when adding an SMB share.



Note: If the **Automatically Create User Directory** setting is selected, for any directory that is automatically created, an ACL is created with the equivalent of UNIX 700 mode bit permissions.

- To maintain the existing permissions on the shared directory, click the **Do not change existing permissions** option.



Note: If the **Automatically Create User Directory** setting is selected, automatically-created directories are assigned UNIX 700 mode bit permissions and no ACL is created.

6. Optionally configure user and group permissions for the share.



Note: Permissions cannot be granted to users accessing files and directories in anonymous mode. For more information, see "SMB user and group configuration."

7. Optionally click **Show Advanced Parameters** to override the default SMB file sharing settings for this share.

Caution: Changes to these parameters may impact the availability of the SMB file sharing service. You should not make changes to this section unless you have experience working with SMB.

The **Settings** section appears, and displays the following configuration settings:

- File and Directory Permissions**

Setting	Description
Directory create mask	Defines which mask bits are applied when a directory is created.
Directory create mode	Defines which mode bits are applied when a directory is created.
File create mask	Defines which mask bits are applied when a file is created.
File create mode	Defines the mode bits are applied when a file is created.
Create permissions	Indicates the default permissions to apply when a file or directory is created.

- Performance Parameters**

Setting	Description
Change notify	Change notify setting. The acceptable values are all , norecurse , or none .
Oplocks	Indicates if an oplock request is allowed.
Strict locking	Indicates if the server checks for and enforces file locks.

- Security Parameters**

Setting	Description
Allow delete read only	Indicates if read-only files can be deleted.
Allow execute always	Indicates if a user with read access to a file can also execute the file.

Setting	Description
Hide dot files	Allows files that begin with a decimal, such as UNIX configuration files, to be hidden.
Impersonate guest	Allows guest access to the share. The acceptable values are always , bad user , and never .
Impersonate user	Allows all file access to be performed as a specific user. This must be a fully qualified user name.
NTFS ACL support	Allows ACLs to be stored and edited from SMB clients.

8. Click **Submit**.

Modify an SMB share

Follow these steps to modify an SMB share.

1. On the **File Sharing** menu, point to **SMB**, and then click **Summary**.
The **SMB Summary** page appears.
2. In the **Shares** section, select **Edit** in the **Actions** column for the share you wish to modify.
The **SMB Edit Share** page appears.
3. Modify the settings as needed.
4. Click **Submit**.

Delete an SMB share

You can delete existing Windows shares that are no longer needed on the cluster.

1. On the **File Sharing** menu, point to **SMB**, and then click **Summary**.
The **SMB > Summary** page appears.
2. In the **Shares** section, select the share you want to delete.
3. Click **Delete**.
The **Confirm Delete** dialog box appears.
4. Click **Yes** to delete the share.

SMB user and group configuration

You can configure the users and groups associated with an SMB share, and view or modify their share-level permissions.

Permissions are defined below:

- **Run As Root:** Gives users or group members root-user permissions on the system, enabling them to create, modify, update, and delete all files. All new files are created with root-user permissions.
- **Full Control:** Allows users or group members to change security permissions on files and folders in the share directory.
- **Change:** Allows write access to all files and folders in the share directory.
- **Read:** Allows read access to all files and folders in the share directory. This is the default share-level permission.

Add a user or group to an SMB share permission

Within each SMB share, you can add new share-level permissions for specific users and groups.

1. On the **File Sharing** menu, point to **SMB**, and then click **Summary**.
The **SMB Summary** page appears.
2. In the **Shares** area, next to the name of the share, click **Edit**.
The **SMB Edit Share** page appears.
3. In the **Users and Groups** area, click **Add**.
The **Add Users or Groups** dialog box appears.

4. To search for **Users**, **Groups**, or both, select the **Search for** check boxes.
5. In the **From this location** list, select the domain in which the users or groups are located.
6. In the **Name** and **Description** lists, select the search parameters **Starts with**, **Exactly matches**, and **Contains**, as needed.
7. Type the **Username** and **Password** for that share.
8. Type the **Name** and **Description** of the users or groups in the boxes.
9. Click **Search**.
The search results appear at the bottom of the **Add Users or Groups** dialog box.
10. Select a user or group from the results list, and then click **Choose**.
The **SMB Edit Share** page appears.

Modify user and group permissions on an SMB share

You can modify user and group permissions for accessing specific shares on the cluster.



Important: It is good practice to only give users the minimum access rights required. Be aware of the following considerations before giving a user "run as root" or "full control" permissions:

- Users or groups that have **Run As Root** permissions will be mapped to "root" when they connect to a share. This setting gives the users or groups unrestricted access to all files and directories on the share. This setting overrides any NTFS access control rules that are configured for files and folders contained in the share.
- Granting a user or group **Full Control** permits unrestricted access to the share. Enable this option only if you want to allow broad access.

1. On the **File Sharing** menu, point to **SMB**, and then click **Summary**.
The **SMB Summary** page appears.
2. Next to the share for which you are modifying user or group permissions, click **Edit**.
The **SMB Edit Share** page appears.



Note: Alternatively, you can modify user or group permissions immediately after adding a share.

3. In the **Users and Groups** area, specify the users or groups whose permissions you want to modify.



Note: This area does not appear if the cluster is in anonymous mode.

- To modify one user or group, click **Edit permissions** next to the name of the share that you want to modify.
- To modify more than one user or group simultaneously, select the check box next to each account name, and then click **Edit selected**.

A permissions dialog box appears.

4. Configure permissions for the user or group.
 - To configure **Run As Root** permissions for the user or group, select the **Run As Root** check box. The check boxes in the **Allow** column are automatically selected.
 - To configure **Full Control**, **Change**, and **Read** permissions individually, select the check boxes in the **Allow** and **Deny** columns as needed.
5. Click **OK**.
The **SMB Edit Share** page appears. The permissions for the user or group are updated.
6. Click **Submit**.
This step is required in order for the changes to take effect.

Remove a user or group from an SMB share permission

You can remove a user or group from an SMB share permission if it is no longer needed.

1. On the **File Sharing** menu, point to **SMB**, and then click **Summary**.

The SMB **Summary** page appears.

2. In the list of shares, click the name of the share that you want to modify.
The SMB **Edit Share** page appears.
3. Click **Edit** next to the share you want to modify.
4. In the **Users and Groups** area, select the name of the user or group permissions that you want to delete, and then click **Delete selected**.



Important: If the cluster is in anonymous mode, the **Users and Groups** area does not appear.

The **Confirm** dialog box appears.

5. Click **Yes**.
The user or group is removed from the list.
6. Click **Submit**.
This step is required in order for the changes to take effect.

NFS (UNIX file sharing protocol)

OneFS supports versions 2 through 4 of the Network File System protocol (NFSv2, NFSv3, NFSv4). You can configure NFS to allow clients to access content stored on Isilon clusters. The Isilon cluster includes a configurable NFS service that enables you to create and manage as many NFS exports as needed.

NFS is enabled by default in the cluster, but can be disabled if necessary. The default `/ifs` export is configured to allow clients to mount any subdirectory, which gives end users access without requiring much administration. You can apply individual host rules to each export, or you can specify all hosts, which eliminates the need to create multiple rules for the same host.

The Isilon cluster supports both asynchronous and synchronous communication over NFS globally.

The Isilon cluster supports the following modes for authentication through NFS:

- **NIS:** Network Information Service (NIS) is a client/server directory service protocol for distributing system configuration data such as user and host names between computers on a network.
- **LDAP:** Lightweight Directory Access Protocol (LDAP) is an application protocol for querying and modifying directory services running over TCP/IP.



Important: You must define all mount points for a given export host as a single export rule. If you want to add an additional mount point for an export host that appears in the list of existing export rules, you must modify that entry rather than add a new one. This holds true for the default host, named `*` (asterisk), as well.

Correct

```
/ifs/home1 /ifs/home2 *
```

Incorrect

```
/ifs/home1 *
/ifs/home2 *
```

View NFS summary information

You can view the status of the NFS file sharing service and manage any exports that have been created on the cluster.

On the **File Sharing** menu, point to **NFS**, and then click **Summary**.

The NFS **Summary** page appears, and displays the following sections:

- **NFS Service:** Displays the status of the NFS file sharing service.
- **Exports:** Displays existing NFS exports and provides the ability to modify, delete, and add new exports to the cluster.

Configure NFS settings

You can enable or disable the NFS service, reload a configuration after making DNS or NIS changes, and set the lock protection level.

1. On the **File Sharing** menu, point to **NFS**, and then click **Settings**.
The **NFS Settings** page appears.
2. View or change the settings in the following sections:
 - **NFS Service Settings:** Displays options for the NFS file sharing service.

Setting	Description
Service	Enables or disables the NFS service.
NFSv4 support	Enables or disables support for NFSv4.
NFSv4 domain	Specifies an NFSv4 domain to be added to local users/groups when sending and receiving NFSv4 principals (user@domain) with NFSv4 clients. The server and clients must agree on this domain name.

- **Exports Settings:** Enables you to reload NFS exports configuration.

Setting	Description
Reload NFS exports configuration	Click Reload to reload NFS exports configuration after making DNS or NIS changes.

- **General Settings:** Allows you to specify the number of nodes that locks will be replicated to.

Setting	Description
Lock protection level	The lock protection level determines the number of node failures that can happen before a lock may be lost.

3. Click **Submit** or **Cancel**.



Note: Changes will not be committed unless you click **Submit**. Clicking **Cancel** will exit the **NFS Settings** page without committing any changes.

Add an NFS export

NFS exports enable UNIX client access to files and directories on the cluster. You can create and manage as many NFS exports as needed.

1. On the **File Sharing** menu, point to **NFS**, and then click **Add export**.
The **NFS Add Export** page appears.
2. In the **Description** box, type a descriptive comment about the export. This setting is not required, but a comment can be useful later if you are managing multiple exports.
3. In the **Directories** box, click **Browse** to locate a directory to export. You can repeat this step to add multiple directory paths.



Note: To prevent problems when setting up new exports, delete any export rules for nonexistent directories, or directories that have been removed from the file system.

4. In the **Clients** box, type the names of the clients that are allowed to access the specified directories. To export all clients, leave this box blank. You can specify a client by host name, IP address, subnet, or netgroup. You can include multiple clients in this box, in the form of one client entry per line.
5. Specify **Permissions** settings:
 - To enable read/write access to the specified directories, select the **Enable write access** check box.
 - To enable read-only access, clear the **Enable write access** check box.
 - To enable mount access to all subdirectories, select the **Enable mount access to subdirectories** check box.
6. Select the **Credential mapping** settings for the export. These settings determine the user and group identity on the share, and control access to the export by mapping only root users or all users to a specific user ID and, optionally, group ID.
 - To map only root users to a specific user ID or group ID, click **Map root users** in the list. This setting changes the user name and group membership for only users logged in as "root" on the client.
 - To map all users to a specific user ID or group ID, click **Map all users** in the list. This setting changes the user name and group membership for all users.
7. In the **User name** box, type the user name to which the specified users will be mapped, or type **nobody** to leave the specified users unmapped.
8. In the **Group membership** box, specify the type of group mapping to perform for the specified users:
 - To retain the current group mapping for specified users, click **Don't modify**.
 - To unmap the specified users from any group membership, click **Set to nobody**.
 - To map the specified users to a specific group, click **Assign groups** and then, in the **Group names** box, type a colon-separated list of group names.
9. Optionally click **Show advanced settings** to override the default NFS file sharing settings for this share.



Caution: Changes to these settings may impact the availability of the NFS file sharing service. You should not make changes to this section unless you have experience working with NFS.

The **Advanced Settings** section appears, and displays the following configurable settings:

- **Performance Settings**

Setting	Description
Commit asynchronous	If set to yes , allows NFSv3 and NFSv4 COMMITs to be asynchronous. The default setting is no .
Directory read transfer preferred	Preferred directory read transfer size reported to NFSv3 and NFSv4 clients.
Read transfer max	Maximum read transfer size reported to NFSv3 and NFSv4 clients.
Read transfer multiple	Recommended read transfer size multiple reported to NFSv3 and NFSv4 clients.
Read transfer preferred	Preferred read transfer size reported to NFSv3 and NFSv4 clients.
Readdirplus prefetch	Number of file nodes to be prefetched on readdir.
Setattr asynchronous	If set to yes , performs set attribute operations asynchronously. The default setting is no .
Write data sync action	Action to perform for DATASYNC writes.
Write data sync reply	Reply to send for DATASYNC writes.
Write file sync action	Action to perform for FILESYNC writes.

Setting	Description
Write file sync reply	Reply to send for FILESYNC writes.
Write transfer max	Maximum write transfer size reported to NFSv3 and NFSv4 clients.
Write transfer multiple	Recommended write transfer size multiple reported to NFSv3 and NFSv4 clients.
Write transfer preferred	Preferred write transfer size reported to NFSv3 and NFSv4 clients.
Write unstable action	Action to perform for UNSTABLE writes.
Write unstable reply	Reply to send for UNSTABLE writes.

- **Access Settings**

Setting	Description
Security flavors	Supported security flavors. Select any or all of the following values: Unix (sys) , Kerberos5 , Kerberos5 integrity , Kerberos5 privacy .

- **Client Compatibility Settings**

Setting	Description
Max file size	Maximum file size.
Readdirplus enable	If set to yes , enables readdirplus. This is the default setting.
Return 32 bit file IDs	If set to yes , returns 32-bit file IDs. The default setting is no .

- **Export Behavior Settings**

Setting	Description
Encoding	Overrides the cluster's general encoding settings for this export.

10. Click **Submit**.

Modify an NFS export

Perform the following steps to modify an NFS export.

1. On the **File Sharing** menu, click **NFS**, and then click **Summary**.
The NFS **Summary** page appears.
2. Click **Edit** next to the NFS export that you want to modify.
The NFS **Edit Export** page appears.
3. Modify the settings as needed.
4. Click **Submit**.

Delete an NFS export

Perform the following steps to delete an NFS export.

1. On the **File Sharing** menu, click **NFS**, and then click **Summary**.

The NFS **Summary** page appears.

2. In the list of exports, click to select the export that you want to delete.
3. Click **Delete**.
The **Confirm Delete** dialog box appears.
4. Click **Yes**.

HTTP and DAV

The Isilon cluster comes equipped with a configurable HTTP service. The cluster uses the HTTP service in two ways: as a means to request files stored on the cluster, and to interact with the web administration interface. The cluster also provides for DAV services, which enable multiple users to manage and modify files.

The DAV (Distributed Authoring and Versioning) service enables multiple users to manage and modify files. DAV is a set of extensions to HTTP that allows clients to read and write from the cluster through the HTTP protocol. You can enable DAV in the web administration interface.

Each node in the Isilon cluster runs an instance of the Apache HTTP Server to provide HTTP access. You can configure the HTTP service to run in one of three modes:

- **Enabled:** Allows HTTP access for cluster administration and browsing content on the cluster.
- **Disabled:** Redirects traffic to the web administration interface. Allows only administrative access to the web administration interface. This is the default mode.
- **Disabled entirely:** Closes the HTTP port used for file access. Users can still access the web administration interface, but they must specify the port number (8080) in the URL in order to do so.

The Isilon cluster also supports a form of the web-based Distributed Authoring and Versioning (WebDAV) protocol that enables users to modify and manage files on remote web servers. The Isilon cluster performs distributed authoring, but does not support versioning and does not perform any security checking.

Configure HTTP and DAV

You can configure HTTP and DAV to enable users to edit and manage files collaboratively across remote web servers.

1. On the **File Sharing** menu, click **HTTP**.
The **HTTP** page appears.
2. Next to the **Service** option, select an access option as needed.
 - To open the HTTP channel for file access and administrative purposes, click **Enable HTTP**.
 - To open the HTTP channel for administrative purposes while disallowing file access, click **Disable HTTP and redirect to the web interface**.
 - To close the HTTP port used for file access, click **Disable HTTP entirely**.



Note: Clicking **Disable HTTP entirely** closes the HTTP port for file access and also disables the redirect; however, you can access the web administration interface at `https://<ip_address>:<port>`, where `<ip_address>` is the IP address of any node in the cluster and `<port>` is the port number for the administration interface (8080, by default).

3. In the **Document root directory** box, type the full path beginning at `/ifs` or click **Browse** to navigate to an existing directory, or click the **File System Explorer** link to create a new directory and set its permissions.



Note: The HTTP server runs as the daemon user and group. For access controls to be properly enforced, the daemon user or group must be granted read access to all files under the document root. The HTTP server must also be able to traverse the document root.

4. In the **Server hostname** box, type the HTTP server name.

The server hostname must be a fully-qualified, SmartConnect zone name and valid DNS name. Valid names must begin with a letter and contain only letters, numbers, and hyphens.

5. In the **Administrator email address** box, optionally type an email address to be displayed to users as the primary contact for issues encountered while serving files.
6. To enable HTTP authentication, select one of the following options from the **ADS Authentication** dropdown menu. To disable HTTP authentication, select **Off**.



Note: For integrated authentication, the cluster must be joined to an Active Directory domain. Access controls are enforced on top of the permissions required by the HTTP server; for settings that include access control, both the HTTP server and the authentication user require read access to the file.

- **Basic Authentication Only:** Enables HTTP basic authentication; user credentials are sent in plain text.
 - **Integrated Authentication Only:** Enables HTTP authentication via NTLM, Kerberos, or both.
 - **Integrated and Basic Authentication:** Enables HTTP basic authentication and integrated authentication.
 - **Basic Authentication with Access Controls:** Enables HTTP basic authentication, and enables the Apache web server to perform access checks.
 - **Integrated Authentication with Access Controls:** Enables HTTP authentication via NTLM and Kerberos, and enables the Apache web server to perform access checks.
 - **Integrated and Basic Auth with Access Controls:** Enables HTTP basic authentication and integrated authentication, and enables access checks via the Apache web server.
7. To allow multiple users to manage and modify files, check the **Enable DAV** box.
A set of extensions to the HTTP protocol, Web-based Distributed Authoring and Versioning (WebDAV) enables users to edit and manage files collaboratively across remote web servers.
 8. To disable access logging, check the **Disable access logging** box.



Important: Heavy use of HTTP file sharing produces performance problems when access logging is enabled. Isilon recommends disabling access logging for high-load applications in a trusted environment.

9. Click **Submit**.

FTP

You can configure the File Transfer Protocol (FTP) to upload and download files stored on an Isilon IQ cluster. The Isilon cluster includes a secure FTP (sFTP) service that you can enable for this purpose.

The Isilon cluster supports FTP access. Any node in the cluster can respond to FTP requests through any standard user account. By default, the FTP service is disabled. When configuring FTP access, ensure that the specified FTP root is the home directory of the user who logs in. For example, the FTP root for local user jsmith would be `/ifs/home/jsmith`. You can enable the transfer of files between remote FTP servers, and you can enable anonymous FTP service on the root by creating a local user named anonymous or ftp.

Enable FTP

Perform the following steps to enable FTP.

1. On the **File Sharing** menu, click **FTP**.
The **FTP** page appears. In the **Service** area, **Disable** is selected by default.
2. Click **Enable**, and then click **Submit**.

Configure FTP settings

You can configure FTP settings to enable server-to-server file transfers, and anonymous or local access to files and directories through the File Transfer Protocol (FTP).

1. On the **File Sharing** menu, click **FTP**.
The **FTP** page appears.
2. In the **Service** area, click **Enable**.
The FTP service is disabled by default.

3. In the **Settings** area, enable or disable FTP settings as needed.

- To enable the transfer of files between two remote FTP servers, next to **Server-to-server transfers**, click **Enable**. This setting is disabled by default.
- To allow users to access files and directories with "anonymous" or "ftp" as the user name and any email address as the password, next to **Anonymous access**, click **Enable**. This setting is disabled by default. With this setting enabled, authentication is not required.



Note: Enabling this setting gives users restricted access to public files, which they can download or upload within a defined area.

- To allow local users to access files and directories with their local user name and password, next to **Local access**, click **Enable**.



Note: Enabling this setting allows local users to upload files directly through the file system. This setting is enabled by default.

4. Click **Submit**.

OneFS data protection

The Isilon cluster is designed to continuously serve data, even when one or more components simultaneously fail. Data protection is applied at the file level, not the block level, enabling the system to recover data quickly. Metadata and inodes are protected at the same level of protection as the data they reference.

Because all data, metadata, and parity information is distributed across all nodes in the cluster, the Isilon cluster does not require a dedicated parity node or drive. This ensures that no single node limits the speed of the rebuild process.

The Isilon system provides several levels of configurable data protection settings, which you can modify at any time without taking the cluster or file system offline or needing to reboot.

With N+1 protection, data is fully available in the event of a failed disk, multiple failed disks within a node, or the failure of an entire node. With N+2, N+2:1, N+3, N+3:1, and N+4 protection, Isilon offers high levels of protection while also taking advantage of parity. The configured protection level affects failure tolerance, with tolerance of up to four failures at one time on distinct nodes. When planning your storage solution, keep in mind that increasing the parity protection settings can affect write performance and requires additional storage space for the increased number of nodes. Although the storage pool is increasing, the reliability scales along with it due to the speed with which a drive can be repaired. As an example, a 250 GB drive can be rebuilt in as little as one hour.

Data layout and file striping

OneFS uses the backend network to automatically allocate and stripe data across the cluster. As the system writes the data, it is also protecting the data at the specified level. No separate action is necessary to protect the data.

The system breaks files into smaller logical chunks called *stripes* before writing them to disk; the size of each file chunk is referred to as the stripe unit size. Each OneFS block is 8 KB, and a stripe unit consists of 16 blocks, for a total of 128 KB per stripe unit. During a write, the system breaks data into stripes and then logically places the data in a stripe unit. As the system lays data across the cluster, it fills the stripe units until the maximum width of the cluster is reached and a parity stripe unit is created. OneFS stripes data across nodes and disks. The stripe width is determined by the number of nodes and the protection setting (for example, N+2).

The system can continuously reallocate data and make storage space more usable and efficient. Depending on the file size and the stripe width (determined by the number of nodes), as the cluster size increases, the system stores large files more efficiently. Due to the way that OneFS applies parity, in a cluster with N+1 protection, files that are 128 KB or smaller are mirrored; in a cluster with N+2 protection, files are triple-mirrored; and so on.

Isilon FlexProtect

The Isilon clustered storage system provides a proprietary system called FlexProtect, which detects and repairs files and directories that are in a degraded state.

Isilon FlexProtect protects data in the cluster based on the configured protection policy, quickly rebuilding failed disks, harnessing free storage space across the entire cluster to further prevent data loss, and monitoring and preemptively migrating data off of at-risk components.

FlexProtect distributes all data and error-correction information across the entire Isilon cluster and ensures that all data remains intact and accessible even in the event of simultaneous component failures.

The supported FlexProtect data protection levels are:

- **N+1:** The cluster can absorb the failure of any single drive or the unscheduled shutdown of any single node without causing any loss in stored data.
- **N+2:1:** The cluster can recover from two simultaneous drive failures or one node failure without sustaining any data loss.
- **N+2:** The cluster can recover from two simultaneous drive or node failures without sustaining any data loss.
- **N+3:1:** The cluster can recover from three simultaneous drive failures or one node failure without sustaining any data loss.
- **N+3:** The cluster can recover from three simultaneous drive or node failures without sustaining any data loss.
- **N+4:** The cluster can recover from four simultaneous drive or node failures without sustaining any data loss.



Note: In terms of overall cluster performance and resources, using a global protection setting of N+1, N+2:1, N+2, N+3:1, N+3, or N+4 is more efficient than using a mirrored protection setting.

N+M data protection

The Isilon system uses the Reed Solomon algorithm for N+M protection.

In the N+M data protection model, N represents the number of nodes, and M represents the number of simultaneous failures of nodes or drives—or a combination of nodes and drives—that the cluster can withstand without incurring data loss. N must be larger than M. OneFS supports N+1, N+2:1, N+2, N+3:1, N+3, and N+4 data protection schemes, and up to 8x mirroring.

For most nodes, the default protection policy is N+1, which means that one drive, multiple drives within a node, or an entire node can fail without causing any data loss. Optionally, you can enable N+2, N+3, or N+4 protection, which allows the cluster to sustain two, three, or four simultaneous failures without causing data loss. You can also optionally enable N+2:1 data protection, which allows the cluster to sustain the loss of two drives or one node without data loss, or N+3:1 data protection, which allows the cluster to sustain the loss of three drives or one node without data loss. Nodes larger than 18 TB default to N+2:1.



Note: For 4U Isilon IQ X-Series and NL-Series nodes, and IQ 12000x/EX 12000 combination platforms, the minimum cluster size of three nodes requires a minimum protection level of N+2:1.

For each protection level, the cluster must contain a minimum number of nodes:

Data protection level	Minimum number of nodes required
N+1	3 nodes
N+2:1	3 nodes
N+2	5 nodes
N+3:1	3 nodes
N+3	7 nodes

Data protection level	Minimum number of nodes required
N+4	9 nodes

OneFS enables you to modify the protection policy in real time, while clients are attached and are reading and writing data. Note that increasing a cluster's protection level will increase the amount of space consumed by the data on the cluster.

The parity overhead for each protection level depends on the file size and the number of nodes in the cluster. The percentage of parity overhead declines as the cluster gets larger. In general, +1 protection has a parity overhead equal to one node's capacity, +2 protection has a parity overhead equal to two nodes' capacity, +3 is equal to three nodes' capacity, and so on.

Number of nodes	+1 overhead	+2:1 overhead	+2 overhead	+3:1 overhead	+3 overhead	+4 overhead
3 nodes	2+1 (33%)	4+2 (33%)	3x	3+3 (50%)	3x	3x
4 nodes	3+1 (25%)	6+2 (25%)	2+2 (50%)	9+3 (25%)	4x	4x
5 nodes	4+1 (20%)	8+2 (20%)	3+2 (40%)	12+3 (20%)	4x	5x
6 nodes	5+1 (17%)	10+2 (17%)	4+2 (34%)	15+3 (17%)	3+3 (50%)	5x
7 nodes	6+1 (14%)	12+2 (14%)	5+2 (28%)	16+3 (15%)	4+3 (43%)	5x
8 nodes	7+1 (12.5%)	14+2 (12.5%)	6+2 (25%)	16+3 (15%)	5+3 (38%)	4+4 (50%)
9 nodes	8+1 (11%)	16+2 (11%)	7+2 (22%)	16+3 (15%)	6+3 (33%)	5+4 (44%)
10 nodes	10+1 (10%)	16+2 (11%)	8+2 (20%)	16+3 (15%)	7+3 (30%)	6+4 (40%)
12 nodes	11+1 (9%)	16+2 (11%)	10+2 (17%)	16+3 (15%)	9+3 (25%)	8+4 (33%)
14 nodes	13+1 (8%)	16+2 (11%)	12+2 (15%)	16+3 (15%)	11+3 (21%)	10+4 (29%)
16 nodes	15+1 (6%)	16+2 (11%)	14+2 (13%)	16+3 (15%)	13+3 (19%)	12+4 (25%)
18 nodes	16+1 (5%)	16+2 (11%)	16+2 (11%)	16+3 (15%)	15+3 (17%)	14+4 (22%)
20 nodes	16+1 (5%)	16+2 (11%)	16+2 (11%)	16+3 (15%)	16+3 (15%)	16+4 (20%)
30 nodes	16+1 (5%)	16+2 (11%)	16+2 (11%)	16+3 (15%)	16+3 (15%)	16+4 (20%)

Data mirroring

The Isilon system supports optional data mirroring from 2x-8x, allowing from two to eight mirrors of the specified content.

Data mirroring requires significant overhead and may not always be the best data-protection method. For example, if you enable 3x mirroring, the specified content is explicitly duplicated three times on the cluster; depending on the amount of content being mirrored, this can require a significant amount of capacity.

Metadata and inodes

The Isilon cluster dynamically allocates metadata as needed; it does not limit the total number of inodes available, and does not place them at risk in a single point of failure such as a metadata server.

All metadata, including information describing where data is stored, file-protection settings, access-control information, and time stamps, is fully protected.

Protection level management

The Isilon clustered storage system provides cluster-wide FlexProtect data protection, which protects data in the cluster based on the configured protection policy. FlexProtect detects and repairs files and directories that are in a degraded state.

The FlexProtect job is responsible for reprotecting data in the event of a node failure.

You can apply protection settings globally across the cluster, or at the file or directory level. This flexibility enables you to protect disparate sets of data at different levels. Note that any child object inherits its policy from the parent, so it is important to ensure that any parent object has an equal or greater level of protection than any of its subordinate objects. For example, if a child directory were protected at the N+2 level, it could tolerate a double device failure; however, if the parent directory were protected at the N+1 level, it could sustain only one loss.

If you enable FlexProtect level +1, +2:1, +2, +3:1, +3, or +4, OneFS invokes the SetProtection phase of the restriper and automatically applies the specified data protection level to all data in the cluster. In most cluster deployments, applying a global FlexProtect setting provides sufficient data protection. However, if necessary, you can manually modify data protection settings at the cluster, directory, or file level by enabling the FlexProtect Advanced option.

Note that OneFS allows you to specify a protection level that the cluster is currently incapable of matching. If you specify an unmatchable protection level, the cluster will continue trying to match the requested protection level until a match is possible. For example, in a four-node cluster, you might specify a 5x protection level; in this example, OneFS would protect the data at 4x until you added a fifth node to the cluster, at which point OneFS would reprotect the data at the 5x level.

Data backup

Isilon supports the Network Data Management Protocol (NDMP) versions 3 and 4. NDMP is an open-standard protocol that provides interoperability with leading data-backup products.

Isilon supports:

- Full backups using NDMP v3 or NDMP v4
- Full restores using NDMP v3 or NDMP v4
- Direct access restore (DAR), single-file restores, and three-way backups
- Incremental backups
- Restore-to-arbitrary systems
- Integration with access control lists (ACLs) such as Posix and NT
- Integration with alternate data streams (NTFS)

You can configure and manage data backups, and view backup log output and backup-job statistics, in the web administration interface. Note that NDMP is disabled by default.

You can optionally add to your cluster one or more Isilon IQ Backup Accelerator nodes to improve backup-and-restore performance. The Backup Accelerator node provides enhanced high-speed backup of file-system data to locally attached tape or media-changer devices that are connected through Fibre Channel on the back end. The Backup Accelerator node supports a wide range of data management applications (DMAs), tape libraries, and switches.

Smart failure and recovery

OneFS supports the ability to smartfail components, reducing the window of risk by protecting content before the device completely fails.

Under normal operating conditions, all data on the cluster is protected against one or more failures of a node or drive. If a node or drive fails, the protection status is considered to be in a degraded state until the data has been reprotected to the configured protection level. Being in a degraded state does not mean that data is lost; rather, it means that the data needs to be reprotected.

In the event of a failure, FlexProtect is responsible for restriping and reprotecting content. Smartfailing a device puts the device into quarantine; it is accessed only as a last resort, and then only for read-only operations. First the data on the node or drive is reprotected, and then the node or drive is removed from the cluster. The node can then be replaced.

The cluster does not require a dedicated hot-spare node or drive in order to recover from a component failure. Rebuilding data from nodes and drives is performed in the free space of the cluster. Because a certain amount of free space is required in order to rebuild data after a failure, it is good practice to ensure that the cluster has adequate free space at all times.

As a cluster grows larger, data restriping operations become faster. As you add more nodes, the cluster gains more CPU, memory, and disks that it can use during recovery operations.

Drive failures

If OneFS determines that a drive's health is suspect, the system automatically smartfails the drive and FlexProtect starts rebuilding data to the free space in the cluster.

OneFS preemptively monitors the health of all disk drives. A drive has three possible states:

- **Healthy:** The drive is in its normal operating condition.
- **Smartfailed:** A restripe process is taking place.
- **Not in use:** The drive is either physically not present or has logically been removed.

Occasionally a drive may fail without the system preemptively detecting a problem; in this case, FlexProtect automatically starts rebuilding the data to available free space on the cluster.

After FlexProtect finishes reprotecting data following a smartfail operation, the drive is logically removed from the cluster and the cluster is no longer in a degraded protection mode.

After you confirm that the FlexProtect operation has completed with no errors, you can hot-swap the drive and then add the new drive to the cluster using the web administration interface or the command-line interface.

Node failures

Because node loss is often a temporary issue, FlexProtect does not automatically start reprotecting data when a node fails or goes offline. If a node reboots, the file system does not need to be rebuilt because it remained intact during the temporary failure.

In a +1 configuration, if one node fails, all data is still accessible from every other node in the cluster. If the node comes back online, it rejoins the cluster automatically without requiring a full rebuild.

To maintain an accurate cluster state, if you physically remove a node from the cluster, you must also logically remove the node from the cluster. After you logically remove a node, it automatically reformats its own drives, and resets itself to the factory default settings. The reset occurs only after OneFS has confirmed that all data has been reprotected. You can logically remove a node using the smartfail process.

During the smartfail process, the node that is to be removed is placed in a read-only state while the cluster performs a FlexProtect process to logically move all data from the affected node. After all data migration is complete, the cluster logically changes its width to the new configuration; at this point, it is safe to physically remove the node.

It is important that you use the smartfail process only when you want to permanently remove a node from the cluster.

It is most efficient to add a replacement node to the cluster before failing the old node, because FlexProtect can immediately use the replacement node to rebuild the failed node's data. If you remove the failed node first, FlexProtect must rebuild the node's data into available space in the cluster, and AutoBalance then transfers the data back to the added replacement node.

WORM (write once, read many)

The OneFS file system offers limited support for WORM (write once, read many) data storage. Files that are committed to a WORM state cannot be edited, overwritten, or renamed. If a file retention date has been configured for the file, it cannot be moved or deleted until the retention date expires.

Before you configure WORM storage on your Isilon cluster, be aware of the following requirements and limitations:

- WORM requires a separate license. For additional information, or to activate WORM for your Isilon clustered storage system, contact your Isilon sales representative.

- The current implementation of WORM is not compliance-ready. If your workflow requires WORM compliance, do not use this feature.

To use WORM, you must first configure an empty directory as a WORM root directory. The root directory, as well as any subdirectories you later create within the WORM root, are then WORM-ready. To commit a file to WORM, you must add it to a WORM-ready directory and then remove its write access.

Create a WORM root directory

Before you can commit files to WORM state, you must create a WORM root directory by marking an empty directory as WORM-ready. Any future subdirectories you create within the WORM root directory are also WORM-ready, capable of storing WORM data.

Before you create a WORM root directory, be aware of the following conditions and requirements:

- Only empty directories can be designated as a WORM root directory.
- All future subdirectories of a WORM root directory inherit its WORM settings.
- A WORM root directory cannot be created as a subdirectory of another WORM-ready directory.
- Hard links cannot cross WORM directory boundaries.
- The path to a WORM file cannot be changed.

WORM configuration is performed at the command line.

1. Open a secure shell (SSH) connection to any node in the cluster and log in using the root account.
2. Create a WORM root directory by running the `isi worm create` command on an empty directory.

For example, to mark an existing directory path of `/ifs/data/worm` as a WORM root directory, you would run the following command:

```
isi worm create -p /ifs/data/worm
```



Note: For full command usage and syntax, type `isi worm create -h` at the command prompt.

You can now write files to the directory and commit them to WORM state. You can optionally modify the WORM root directory to specify a default retention period or override the retention date for WORM files in the directory.

Configure a WORM root directory

You can set a default retention period for WORM files at the WORM root directory level in addition to any retention dates that are configured for individual files. You can also override retention dates for files that are committed to WORM state. In either case, if retention dates are set at both the file level and the directory level, the later date is used for WORM file expiration.



Note: By default, a WORM file is set to expire immediately upon committing it to WORM state.

1. Open a secure shell (SSH) connection to any node in the cluster and log in using the root account.
2. To set a default retention period, or offset, for WORM files committed to the WORM root, run the `isi worm modify` command with the `-d` option.

For example, to set a default retention period of 6 weeks from the time a file is committed to the WORM-ready directory `ifs/data/worm`, you would run the following command:

```
isi worm modify -p /ifs/data/worm -d 6w
```



Note: For command usage and syntax, type `isi worm modify -h`.

The default retention offset is applied only to WORM files that do not have a retention date configured.

3. To override the retention date for files in a WORM-ready directory, run the `isi worm modify` command with the `-o` option.

For example, to set a retention date of December 22, 2012 for files that are committed to the `ifs/data/worm` directory, you would run the following command:

```
isi worm modify -p /ifs/data/worm -o 20121222
```



Note: For command usage and syntax, type `isi worm modify -h`.

Files that are committed to the WORM root directory or one of its subdirectories will use the file retention date (if configured) or the override retention date, whichever is later.

Set a WORM file retention date

By default, a WORM file is set to expire immediately when it is committed to a WORM state. An expired WORM file can be moved or deleted, but cannot be otherwise modified while in a WORM-ready directory. You can manually set a WORM file to expire at a later time by configuring a WORM file retention date. If an override retention date is also configured for the WORM root directory, the later date is used for WORM file expiration.



Note: You cannot use SMB to set file retention dates.

1. Open a connection to the cluster via NFS or SSH.
2. Run the `touch -at` command on the file.

For example, to set a retention date of midnight October 4, 2017 for a file named `readme.txt` in the WORM-ready directory `/ifs/data/worm/`, run the following command:

```
touch -at 201710040000 /ifs/data/worm/readme.txt
```



Note: For command syntax and usage, type `touch -h`.

Commit a file to WORM state

After a directory has been marked as WORM-ready, files within the directory can be committed to a WORM state through SMB, NFS, or local operations. While in WORM state, a file cannot be moved or altered in any way; however, a WORM file can be forcibly deleted.



Note: Prior to committing a file to WORM state, you can set a retention date at which the WORM state will expire. For more information, see "Set a file retention date."

1. Open a connection to the cluster via SSH, NFS, or an SMB mount.
2. Write a file to a WORM-ready directory.
3. Remove write access to the file.
 - To remove write access by using the command line, run the `chmod ugo-w` command on the file. For command usage and syntax, type `chmod -h`.
 - To remove write access by using SMB, set the following parameter:
`read only = yes`
 - To remove write access by using the web administration interface, follow these steps:
 1. On the **File System** menu, click **File System Explorer**.
 2. Navigate to the file, and then click **Properties**.
 3. In the UNIX Permissions section, clear each check box in the **Write** column.

List WORM root directories

You can use the `isi worm list` command to list WORM root directories on the cluster.

At the command prompt, run the `isi worm list` command. (For command usage and syntax, type `isi worm list -h`.)

- To view basic information about each WORM root on the cluster, type:

```
isi worm list
```

If a WORM root directory exists, the system displays output similar to the following example:

Root Path	Type
/ifs/data/worm	WORM

- To view more detailed information including the WORM domain ID, override retention date, and default minimum retention offset, use the `-l` option:

```
isi worm list -l
```

If a WORM root directory exists, the system displays output similar to the following example:

ID	Root Path	Type	Override Ret. Date	Default Min. Ret.
65537	/ifs/data/worm	WORM	2012-12-22 00:00:00	2y

View WORM file and directory information

You can use the `isi worm info` command to view WORM information for a specific file or directory.

At the command prompt, run the `isi worm info` command on a file or directory. You must specify the absolute path using the `-p` option.



Note: For command usage and syntax, type `isi worm info -h` at the command prompt.

For example, for the following scenarios assume that `/ifs/data/worm` is a WORM-ready directory containing WORM and non-WORM files.

- To view WORM information about the `/ifs/data/worm` directory, type:

```
isi worm info -p /ifs/data/worm/
```

The system displays output similar to the following example:

```
Domain Info for: /ifs/data/worm

Domains
-----
ID      | Root Path      | Type
-----+-----+-----
65537   | /ifs/data/worm | WORM
```

This output indicates that the directory is a WORM root whose ID is 65537.

- To view WORM information about the `/ifs/data/` directory, type:

```
isi worm info -p /ifs/data/
```

The system displays output similar to the following example:

```
Domain Info for: /ifs/data

Domains
-----
None

Contains Directories with These WORM Domains
-----
ID      | Root Path      | Type
-----+-----+-----
65537   | /ifs/data/worm | WORM
```

This output indicates that the `/ifs/data/` directory is the parent of a WORM root directory.

- To view WORM information about a file named `file.txt` in the `/ifs/data/worm/wormsub` directory, type:

```
isi worm info -p /ifs/data/worm/wormsub/file.txt
```

The system displays output similar to the following example:

```
Domain Info for: /ifs/data/worm/wormsub/file.txt

Domains
-----
ID      | Root Path      | Type
-----+-----+-----
65537   | /ifs/data/worm | WORM

Worm State: not committed   Would Expire: Never
```

In this example, `file.txt` is in a WORM directory tree rooted at `/ifs/data/worm`, but has not been committed to a WORM state.

Delete a WORM file

You can use the `isi worm filedelete` command to forcibly delete a file, regardless of its WORM state.

1. Open a secure shell (SSH) connection to any node in the cluster and log in using the root account.
2. Run the following command, where `<Path-to-WormFile>` is the absolute path to the file that you want to delete:

```
isi worm filedelete -p <Path-to-WormFile>
```

For example, to delete a WORM file named `readme.txt` from the `/ifs/data/worm/` directory, type:

```
isi worm filedelete -p /ifs/data/worm/readme.txt
```

File System Explorer

File System Explorer is a web-based interface that enables you to manage the content stored on the cluster. You can use File System Explorer to navigate the Isilon file system, add directories, and manage file and directory properties including data protection, I/O optimization, and UNIX permissions.



Important: Isilon file system (`/ifs`) directory permissions are initially set to allow full access for all users--a setting that is convenient but not secure. Any user can delete any file, regardless of the permissions on the individual file. Depending on your particular environment, this might or might not be acceptable. You may need to change the default configuration by establishing some basic permission restrictions.



Note: You can view and configure file and directory properties from within Windows clients that are connected to the cluster. However, because Windows and UNIX permissions differ from one another, you must be careful not to make any unwanted changes that affect file and directory access.

Navigate the OneFS file structure

You can use the File System Explorer to navigate the Isilon file system and view the contents of individual directories.

The File System Explorer layout is similar to that of most other graphical representations of files and directories, with hierarchies of individual files and directories.

On the **File System** menu, click **File System Explorer**.

The File System Explorer interface appears.



Note: You can adjust the sizes of the directory (left-side) and content (right-side) panes by clicking and dragging the vertical divider to the left or right.

- To view a directory's subdirectories, click the plus sign (+) next to the directory name in the left pane.
- To view a directory's contents including files and subdirectories, click the directory name in the left pane. The contents of the directory appear in the right pane, including basic information about each file and subdirectory.



Note: If the directory contents span multiple pages, you can navigate the results by typing a page number in the **Page** box, or by clicking the right or left arrow icons.

- To search for files and subdirectories within the current directory, type all or part of the name in the text box, and then click **Search**.

View file and directory properties

You can view a file or directory's protection settings, I/O optimization settings, UNIX permissions, location, and last-modified/last-accessed times. Additional file properties include file size and disk space used. Additional directory properties include a content summary, SMB share settings, and NFS export settings.



Note: As an alternative to the method described here, you can view file and directory properties through any Windows clients that are connected to the cluster.

1. On the **File System** menu, click **File System Explorer**.
The File System Explorer interface appears.
2. In the contents pane, under **Actions**, click the **Properties** link for the file or directory whose properties you want to view.
The **Properties** page appears.
3. Review the basic directory or file properties.
 - For directories, review the **Directory name**, the **Path**, the **Contents** (number of files and directories), the dates on which the directory was **Last modified** and **Last accessed**, and whether the file or directory is shared through Windows File Sharing (SMB) or exported through UNIX file sharing (NFS).
 - For files, review the **Filename**, the **Path**, the **File size** and **Space used**, and the dates on which the file was **Last modified** and **Last accessed**.
4. Review the **Protection Settings** section.



Note: If these settings are overridden by SmartPools (as indicated by a notification message), you must navigate to the SmartPools **Settings** page in order to view them.

- **Settings management:** Displays whether these settings are manually managed or managed by SmartPools.
- **Disk pool:** Displays the disk pool whose protection policy will be used if SmartPools is configured to manage protection settings. This setting is only available if SmartPools is licensed and enabled on the cluster.

- **SSD strategy:** Displays the strategy that will be used for user data and metadata if solid-state drives (SSDs) are available. You can choose to use or avoid SSDs for user data, or use *metadata acceleration*, which creates a mirror backup of metadata on an SSD and writes the rest of the metadata plus all user data on hard disk drives (HDDs).
- **Protection level:** Displays the FlexProtect or data mirroring protection policy that is configured for the file or directory. If SmartPools is licensed and enabled on the cluster, by default the file or directory inherits the protection policy that is configured for the specified disk pool.

5. Review the **I/O Optimization Settings** section.



Note: If these settings are overridden by SmartPools (as indicated by a notification message), you must navigate to the SmartPools **Settings** page in order to view them.

- **Settings management:** Displays whether these settings are manually managed or managed by SmartPools.
- **SmartCache:** Displays whether SmartCache is enabled or disabled. SmartCache enables you to accelerate the process of writing content to the cluster.
- **Data access pattern:** Displays the optimization settings for accessing data (**Concurrency**, **Streaming**, or **Random**). By default, iSCSI LUNs are configured to use a random access pattern; other files and directories use a concurrent access pattern by default.

6. Review the **UNIX permissions** section.

- **User:** Displays the name of the owner user of the selected file or directory.
- **Group:** Displays the name of the owner group of the selected file or directory.
- **Permissions:** Displays the read, write, and execute permissions granted to the owner user, members of the owner group, and other users on the system.

Modify protection settings

The OneFS operating system supports several levels of protection for files and directories on the cluster, including Isilon FlexProtect (+1 through +4) and data mirroring (2x-8x).



Note: If these settings are overridden by SmartPools (as indicated by a notification message), you must navigate to the SmartPools **Settings** page in order to modify them.

1. On the **File System** menu, click **File System Explorer**.
The File System Explorer interface appears.
2. In the contents pane, under **Actions**, click the **Properties** link for the file or directory whose properties you want to view.
The **Properties** page appears.
3. In the **Protection Settings** section, modify the following settings as needed.



Note: If you are modifying a directory, you can apply each individual setting to the directory's contents by clicking its associated **Apply setting to contents** check box.

- **Settings management:** Click to specify whether to manage these settings manually or via SmartPools.



Note: If you change either or both settings for the disk pool or protection policy, this field automatically refreshes to display **Manually managed**. If you select **Managed by SmartPools**, the disk pool and protection policy values will refresh to display the applied settings the next time the SmartPools job runs.

- **Disk pool:** Click the name of a disk pool to apply its protection policy settings to this file or directory. This setting is only available if SmartPools is licensed and enabled on the cluster.
- **SSD strategy:** Click to select an SSD strategy to use for this file or directory.

- **Metadata acceleration:** Creates a mirror backup of file metadata on an SSD and writes the rest of the metadata plus all user data on HDDs. Depending on the global namespace acceleration setting, the SSD mirror may be an extra mirror in addition to the number required to satisfy the protection level.
- **Avoid SSDs:** Never uses SSDs; writes all associated file data and metadata to HDDs only.
- **Data on SSDs:** Similar to metadata acceleration, but also writes one copy of the file's user data (if mirrored) or all of the data (if not mirrored) on SSDs. Regardless of whether global namespace acceleration is enabled, any SSD blocks reside on the file's target pool if there is room. This SSD strategy does not result in the creation of additional mirrors beyond the normal protection level.
- **Protection level:** Click to select a FlexProtect or data mirroring protection policy for this file or directory. By default, the file or directory inherits the protection policy that is configured for the specified disk pool.

4. Click **Submit**.

Modify I/O optimization settings

The OneFS operating system supports SmartCache, which can greatly accelerate the process of writing content to the cluster. You can apply SmartCache to a single file or directory, or a directory and all of its contents.



Note: If these settings are overridden by SmartPools (as indicated by a notification message), you must navigate to the SmartPools **Settings** page in order to modify them.

1. On the **File System** menu, click **File System Explorer**.
The File System Explorer interface appears.
2. In the contents pane, under **Actions**, click the **Properties** link for the file or directory whose properties you want to view.
The **Properties** page appears.
3. In the **I/O Optimization Settings** section, modify the following settings as needed.



Note: If you are modifying a directory, you can apply each individual setting to the directory's contents by clicking its associated **Apply setting to contents** check box.

- **Settings management:** Click to specify whether to manage these settings manually or via SmartPools.



Note: If you change either or both settings for SmartCache or data access, this field automatically refreshes to display **Manually managed**. If you select **Managed by SmartPools**, the SmartCache and data access pattern values will refresh to display the applied settings the next time the SmartPools job runs.

- **SmartCache:** Click to specify whether SmartCache is enabled or disabled. This setting is disabled by default.



Note:

SmartCache can accelerate the process of writing content to the cluster. However, writes to the cluster using SMB will not be affected by SmartCache when write through is specified by the client. Also, writes to the cluster using NFS will not be affected by SmartCache if they are tagged as stable.

If a node crashes with SmartCache enabled while using NFS, unstable writes that have not been committed will be temporarily lost on the node. However, the uncommitted data will be sent to the cluster again as soon as the client reconnects to the cluster, unless the client crashes. As long as the client does not crash before it reconnects, no data will be lost.

- **Data access pattern:** Click to change the optimization settings for accessing data. Available options include:
 - **Concurrency:** Optimizes for concurrent load on the cluster, that is, many simultaneous clients. With the exception of iSCSI logical units, this is the default setting.
 - **Streaming:** Optimizes for high-speed streaming of a single file, for example to enable very fast reading with a single client.

- **Random:** Optimizes for unpredictable access to the file, by disabling prefetch. This is the default setting for iSCSI logical units.

4. Click **Submit**.

Modify UNIX permissions

You can modify the user and group ownership of files and directories, and set permissions for the owner user, owner group, and other users on the system.

1. On the **File System** menu, click **File System Explorer**.
The File System Explorer interface appears.
2. In the contents pane, under **Actions**, click the **Properties** link for the file or directory whose properties you want to view.
The **Properties** page appears.
3. In the **UNIX Permissions** section, modify the following settings as needed.



Note: If you are modifying a directory, you can apply each individual setting to the directory's contents by clicking its associated **Apply setting to contents** check box.

- **Owner:** By default, files and directories are owned by the root user account. To specify a different owner, type the name of an existing user account in the text box or click the down arrow to select a user from the list.
- **Group:** By default, files and directories are owned by the wheel group. To change the group ownership, type the name of an existing group in the text box or click the down arrow to select a group from the list.
- **Permissions:** Select or clear the check boxes to assign read/write/execute permissions to the specified account owner (**User**), group members (**Group**), and anyone (**Other**). Refer to the following table for explanations of individual settings.

User or group	Access type	File	Directory
User	Read	The account owner can read the file.	The account owner can list the directory contents.
	Write	The account owner can modify the file.	The account owner can create, delete, and rename files in the directory.
	Execute	The account owner can execute the file.	The account owner can access the directory.
Group	Read	Members of the owner group can read the file.	Members of the owner group can list the directory contents.
	Write	Members of the owner group can modify the file.	Members of the owner group can create, delete, and rename files in the directory.
	Execute	Members of the owner group can execute the file.	Members of the owner group can access the directory.
Other	Read	Anyone can read the file.	Anyone can list the directory contents.
	Write	Anyone can modify the file.	Anyone can create, delete, and rename files in the directory.
	Execute	Anyone can execute the file.	Anyone can access the directory.

4. Click **Submit**.

Add a directory

You can add directories to your file system through the Isilon File System Explorer at any time.

1. On the **File System** menu, click **File System Explorer**.
The File System Explorer interface appears.
2. Click **Add Folder**.
The **New Folder Properties** menu appears.
3. In the **Folder name** box, type a name for the new directory.



Note: By default, the new directory is created within the current directory. You can use forward slashes (/) to specify a subdirectory or the full path beginning at /ifs where you want to create the new directory.

4. From the drop-down lists, select the **Owner** and **Group** for the directory.
5. Select the check boxes to establish the basic permissions for the directory.
6. Click **Submit**.

NDMP backup management

Isilon IQ clustered storage systems support the Network Data Management Protocol (NDMP) open standard, which enables interoperability with several leading backup products.

Isilon clustered storage systems support full backup and restore operations through NDMP v3 or NDMP v4.

You can optionally add to your cluster one or more Isilon IQ Backup Accelerator nodes to improve backup-and-restore performance. The Backup Accelerator node provides enhanced high-speed backup of file-system data to locally attached tape or media-changer devices that are connected through Fibre Channel on the back end. The Backup Accelerator node supports a wide range of data management applications (DMAs), tape libraries, and switches.

For information about purchasing an Isilon IQ Backup Accelerator node, contact your Isilon sales representative.

Configure NDMP backup settings

To enable the cluster to perform backup over NDMP, you must enable the NDMP service and configure other NDMP backup settings according to your storage area network (SAN) environment.

1. On the **File System** menu, point to **Backup**, and then click **NDMP Settings**.
The **NDMP Settings** page appears.
2. Configure the NDMP **Service** setting:
 - To enable the NDMP service, under **Service**, click **Enable**. This is the default setting.
 - To disable the NDMP service, under **Service**, click **Disable**. When the NDMP service is disabled, the cluster ignores any new NDMP requests from clients. Disabling the NDMP service does not affect any NDMP sessions that are already in progress.
3. Configure the basic NDMP settings:
 - a. Under **Settings**, click **Edit settings**.
The **Edit Settings** dialog box appears.
 - b. In the **Port number** box, type the TCP/IP port number on which the NDMP server listens for incoming requests.
The default setting is **10000**.
 - c. In the **DMA vendor** list, click the name of an Isilon-qualified DMA vendor or, if you are using a non-qualified DMA vendor, click **generic**.
 - d. Click **Submit**.

4. Add one or more NDMP administrators with username/password credentials for DMA authentication. You must define at least one NDMP administrator account.
 - a. Under **NDMP Administrators**, click **Add administrator**.
The **Add Administrator** dialog box appears.
 - b. In the **Name** box, type the username portion of the NDMP administrator's credentials.
 - c. In the **Password** box, type the password portion of the NDMP administrator's credentials.
 - d. In the **Confirm password** box, retype the password for the NDMP administrator.
 - e. Click **Submit**.

The new NDMP administrator appears in the **NDMP Administrator** list.

Manage NDMP backup devices

You can manage tape devices or media changers that are locally attached to an Isilon IQ Backup Accelerator node.

If you locally attach a tape device or a media-changer device to your Isilon IQ Backup Accelerator node, the device is detected when you start or restart the Backup Accelerator node or when you rescan the node's Fibre Channel ports to discover devices. The Backup Accelerator node creates a device entry, which is composed of settings detected during device discovery, for each newly discovered tape device or media-changer device; however, you can modify or delete these device entries as needed. This section describes how to manage entries for tape devices or media-changer devices that are locally attached to a Backup Accelerator node.

After you connect a device to the storage area network (SAN), directly connect a device to a Backup Accelerator node, or make changes to the SAN topology, you must rescan the Fibre Channel ports in order for the Backup Accelerator node to discover the devices or detect the updated SAN topology settings.

1. On the **File System** menu, point to **Backup**, and then click **Devices**.
The Backup **Devices** page appears, and displays information about any locally attached tape devices or media-changer devices.



Note: The Backup **Devices** page displays data only if an Isilon IQ Backup Accelerator node is installed in your cluster, and if any tape devices or media-changer devices are locally attached to the Backup Accelerator node through Fibre Channel connections. To filter out inactive-path information and view information about only currently active paths, click **Show only active paths**.

2. Review the device settings:
 - **Name:** Specifies the device's name.
 - **State:** Indicates whether the device is currently in use (**Open**) or not in use (**Closed**).
 - **WWN:** Specifies the Fibre Channel World Wide Name (WWN) that binds the logical device to the physical device.
 - **Product:** Specifies the device's vendor name and model name or number.
 - **Serial Number:** Specifies the device's serial number.
 - **Paths:** Specifies the name of the Backup Accelerator node, and the port number or numbers, to which the device is connected. To view additional information about a node, click the **Node** link.
 - **LUN:** Specifies the device's logical unit number (LUN).
 - **Port ID:** Specifies the Fibre Channel port ID that binds the logical device to the physical device.
 - **WWPN:** Indicates the device's World Wide Port Name (WWPN).
3. Optionally, you can modify a device's name if, for example, you want to assign a more descriptive or user-friendly name to the device.
 - a. Click the device's **Name** link.
The **Rename Device** dialog box appears.
 - b. In the **Device name** box, type a new name for the device.
 - c. Click **Submit**.

4. Optionally, you can delete a device if you no longer need it by clicking the device's **Delete device** link in the **Actions** column.
5. Optionally, you can rescan the Fibre Channel ports in order to discover new devices or detect changes in your SAN topology.
 - a. Click **Discover devices**.
The **Discover devices** dialog box appears.
 - b. In the **Node** list, click the number of the node that contains the port that you want to rescan, or click **All nodes**.
 - c. In the **Ports** list, click the number of the port that you want to rescan, or click **All ports**.
 - d. If you want to update the displayed device information to omit any devices or ports that no longer exist, select the **Delete inaccessible devices** check box.
 - e. Click **Submit**.
The Backup **Devices** page refreshes and displays updated device information.

Manage NDMP backup ports

Perform the following steps to manage the Fibre Channel ports that are used to locally attach tape devices or media changers to a Backup Accelerator node.

Each Isilon IQ Backup Accelerator node can support up to four paths to each locally connected tape device or media-changer device through four Fibre Channel ports. These multiple paths can be used for load-balancing the Fibre Channel traffic.

This section describes how to manage the Fibre Channel ports that are used to locally attach tape devices or media-changer devices to an Isilon IQ Backup Accelerator node.

1. On the **File System** menu, point to **Backup**, and then click **Ports**.

The Backup **Ports** page appears, and displays information about the four Fibre Channel ports on each Backup Accelerator node.



Note: The Backup **Ports** page displays data only if one or more Isilon IQ Backup Accelerator nodes are installed in your cluster.

2. Review the Fibre Channel port data:
 - **Port:** Displays the name of the Backup Accelerator node, and the node's port number. To view additional information about a node, click the **Node** link.
 - **Topology:** Indicates the type of Fibre Channel topology in use in your storage area network (SAN). In a **Point to Point** topology, two devices are connected back to back. In a **Loop** topology, all devices are in a loop, or ring, formation. If the **Auto** setting is enabled, the Fibre Channel host bus adapter (HBA) is configured to automatically detect the SAN topology type.
 - **WWNN:** Indicates the port's World Wide Node Name (WWNN).
 - **WWPN:** Indicates the port's World Wide Port Name (WWPN).
 - **Rate:** Indicates the port's data rate.
 - **Actions:** Indicates whether the node is currently enabled or disabled. If the **Disable** link is displayed, the port is enabled. If the **Enable** link is displayed, the port is disabled.
3. Optionally, modify a port's settings as needed.
These settings control how the Fibre Channel HBA presents itself. Isilon recommends not modifying these settings unless necessary. For more information, contact Isilon Technical Support.
 - a. For the port whose settings you want to modify, click the corresponding **Port** link in the **Port** column.
The **Edit Port** dialog box appears.
 - b. In the **WWNN** box, type the port's WWNN.
 - c. In the **WWPN** box, type the port's WWPN.
 - d. In the **Topology** list, click the type of Fibre Channel topology in use in your SAN (**Point to Point**, **Loop**, or **Auto**).

The default setting, **Auto**, is recommended in most cases. If the **Auto** setting is enabled, the Fibre Channel HBA attempts to automatically detect the SAN topology settings, and configures itself accordingly. However, if the Backup Accelerator node cannot access the device, and the device is directly connected to the node, changing this setting to **Point to Point** might resolve the access problem. Enable the **Loop to Loop** option only if the device is in an arbitrated loop configuration.

- e. In the **Rate** list, click the port's data rate (**Auto**, **1Gb/s**, **2Gb/s**, or **4Gb/s**).

The default setting, **Auto**, is recommended in most cases. If the **Auto** setting is enabled, the Fibre Channel HBA attempts to automatically detect the port's data rate, and configures itself accordingly.

- f. Click **Submit**.

4. Optionally, in the **Actions** column, enable or disable ports as needed.

- **Disable:** Click to disable the port. If you disable a port, all device paths that use the port are also disabled.
- **Enable:** Click to enable the port.

View NDMP backup logs

You can view log files for NDMP operations to track and troubleshoot backup processes.

Log entries can include important information about recent NDMP operations. Each entry includes a description of the NDMP activity and the time the entry was generated.

On the **File System** menu, point to **Backup**, and then click **Logs**.

The **Logs** page appears, and displays recent log entries for all cluster operations, including NDMP events.

View and manage NDMP backup sessions

You can view information about NDMP backup operations, and you can terminate an NDMP session if necessary.

If the NDMP service is enabled, NDMP data management applications (DMAs) can establish communication with the NDMP daemons running on Isilon IQ storage nodes or, if applicable, Isilon IQ Backup Accelerator nodes. All communication between the DMA and the cluster is performed in the context of an NDMP session. You can review status information about these NDMP sessions, and you can terminate a session if necessary. This information is refreshed every 15 seconds.

1. On the **File System** menu, point to **Backup**, and then click **Sessions**.
The NDMP **Sessions** page appears.
2. Review the session data:
 - **Session:** Unique session identification number.
 - **Elapsed:** Elapsed time since the session started.
 - **Transferred:** Amount of data transferred since the session started.
 - **Throughput:** Average session throughput over the past five minutes.
 - **Client/Remote:** IP address of the DMA client host. If a three-way backup and restore process is currently running, the IP address of the remote tape device or media-changer device also appears.
 - **Mover/Data:** Current status of the NDMP protocol state machines.
 - **Operation:** Type of operation in progress (**Backup** or **Restore**).
 - **Source/Destination:** Names of any source and destination directories that are currently in use, if a backup or restore operation is running. The source directory is the backup root directory. The destination directory, which is displayed only if a restore operation is currently running, is the location to which the data is being restored.
 - **Device:** Name of tape or media-changer device.
 - **Mode:** Open mode, if the device is in use.
 - **Actions:** To terminate a session and release any associated device resources, click **Kill**.

Cluster anti-virus scanning

The OneFS operating system enables enterprise-wide scanning for computer viruses and other security threats on your Isilon storage cluster by integrating with third-party scanning services.

Files on an Isilon cluster can be scanned for threats on a scheduled basis, scanned on-demand when files are accessed by end users, or manually scanned at any time. If viruses are detected on the cluster, the infected files can be repaired, quarantined, or truncated to eliminate the threat. The cluster notifies storage administrators when threats are detected using three methods: alerts, near real-time summary information, and historical reports.

Planning an anti-virus scanning deployment is a balancing act between managing security threats and maintaining cluster performance. Scans can be performed cluster-wide or targeted at selected directories and file types that are more prone to malicious intrusions. Administrator-defined anti-virus scanning policies provide flexibility for scheduling when and where the scans take place. Integrating the cluster with multiple ICAP servers can improve performance with automatic load balance and failover.

Isilon clusters interoperate with the following anti-virus servers using the Internet Content Adaptation Protocol (ICAP):

- Symantec Scan Engine
- Trend Micro Interscan Web Security Suite
- McAfee VirusScan Enterprise for Storage

Cluster anti-virus summary

The Anti-Virus **Summary** page provides an overview of anti-virus scanning activity on the cluster.

From the **Summary** page you can view currently running anti-virus scans, recently completed scans, and recently detected threats. You can also view the status of the cluster's anti-virus scanning service, and the status of third-party ICAP scanning servers configured to communicate with the cluster.

View anti-virus summary information

You can view summary information about anti-virus scanning activity on the cluster.

The Anti-Virus **Summary** page displays overview information about the status of the cluster's anti-virus scanning service and third-party ICAP scan servers, as well as the most recent anti-virus scan results and detected threats.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Summary**.
The Anti-Virus **Summary** page appears.
2. View summary anti-virus information about the cluster in the following sections:
 - **Service:** Indicates whether the cluster anti-virus scanning service is enabled or disabled.
 - **ICAP Servers:** Lists all third-party anti-virus servers configured on the cluster and whether they are enabled or disabled.
 - **Currently Running:** Displays any scheduled or manually initiated anti-virus scans that are currently being run against files on the cluster. Pending scans are denoted by a yellow status indicator icon; pausing the mouse pointer over the yellow icon displays the text *Scan is waiting to run*.
 - **Recently Completed:** Shows the results of the 10 most recently completed anti-virus scans on the cluster, including scan success or failure status, scan start time and duration, virus detection results, number of files scanned, root directories included in the scan, and whether global file restrictions were enforced.
 - **Recently Detected Threats:** Shows the 10 most recent virus threats detected on the cluster, including the threat name, file names, and directory locations of the infected files; time of detection; type of remediation action taken against the infection; and the name of the policy that detected the infection.
3. Optionally, you can pause the mouse pointer over any indicator icon in the **Status** column to view a brief description of the status of an anti-virus scan, detected threat, or the status of the cluster anti-virus service and communication with third-party ICAP scan servers.

The cluster anti-virus scanning service

The cluster's anti-virus scanning service controls whether scans are performed on the cluster.

When the scanning service is enabled, anti-virus scans can be run automatically or manually. When the scanning service is disabled, all current scanning is halted and pending scans cannot proceed.

Before you can enable the anti-virus scanning service, at least one third-party ICAP scan server is configured on the cluster.

You can disable the cluster anti-virus scanning service during system maintenance or to improve cluster performance. However, if threat detection and data security are priorities, keep the service enabled.

Enable the cluster anti-virus scanning service

If you disable the cluster anti-virus scanning service for system maintenance or to improve cluster performance, you must re-enable the service in order to resume anti-virus scanning on the cluster. When you initially add an ICAP scan server to the cluster, the anti-virus scanning service is automatically enabled.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Summary**.
The Anti-Virus **Summary** page appears.
2. Under **Service**, click **Enable**.
The service status indicator icon changes from red to green. (The icon may also briefly turn yellow if the service is not running on all nodes in the cluster.)

Disable the cluster anti-virus service

You can disable the cluster anti-virus scanning service if you need to perform system maintenance or otherwise want to halt all scanning.

Disabling the service interrupts any currently running scans and prevents scheduled scans from running until the service is re-enabled. Disabling the cluster scanning service also prohibits the start of manual scans by hiding the **Start** command in the **Actions** column on the Anti-Virus **Summary** page and the **Policies** page.



Note: If only one ICAP scan server is configured on the cluster, the anti-virus scanning service is automatically disabled if you delete that sole ICAP server.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Summary**.
The Anti-Virus **Summary** page appears.
2. Under **Service**, click **Disable**.
The service indicator icon changes from green to red.

ICAP scan server configuration

Before you can enable anti-virus scanning on a cluster, you must configure the cluster to communicate with one or more third-party ICAP scan servers.

Configuring multiple scan servers provides for automatic failover in case a scan server stops communicating with the cluster, and provides round-robin load balancing among the servers.

The cluster anti-virus server supports the following ICAP scan servers:

- Symantec Scan Engine
- Trend Micro Interscan Web Security Suite
- McAfee VirusScan Enterprise for Storage

Color-coded icons on the Anti-Virus **Summary** page indicate the status of an ICAP scan server:

- Green indicates that the ICAP scan server is enabled and communicating with the cluster.
- Red indicates that the ICAP scan server is down or not communicating with the cluster.
- Gray indicates that the ICAP scan server has been disabled from the cluster.

Add an ICAP scan server to a cluster

You must configure the cluster to communicate with at least one third-party ICAP scan servers before you can enable anti-virus scanning on a cluster. This requires providing an IP address for each scan server.

The minimum configuration requires configuring one ICAP scan server, but you may optionally add more. Configuring multiple scan servers provides for automatic failover in case a scan server stops communicating with the cluster, and provides round-robin load balancing among the servers.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Summary**.
The Anti-Virus **Summary** page appears.
2. Under **ICAP Servers**, click **Add server**.
The **Add ICAP Server** dialog box appears.
3. In the **ICAP URL** box, type the ICAP protocol prefix followed by the IP address of the ICAP scan server.
For example: `icap://10.20.100.3`
Optionally, you can append a port number to the URL. Isilon clusters use port 1344 by default for ICAP communications.
For example: `icap://10.20.100.3:1344`
Optionally, you can append a path to the URL.
For example: `icap://10.20.100.3/avscans`
4. Optionally, type a **Description** for the ICAP scan server, such as the product name or vendor.
The description can be up to 50 characters long.
5. Click **Submit**.
6. Optionally, you can continue to add other scan servers by repeating the preceding steps.

After you add an ICAP scan server to a cluster, the server is automatically enabled, provided that the server is communicating with the cluster, as indicated by the green icon in the **Status** column. If the ICAP scan server is not available or responding, the status indicator icon is red.

Enable an ICAP scan server on a cluster

If you disable an ICAP scan server on the cluster—for maintenance or other reasons—you must enable it again in order to resume anti-virus scanning.



Note: After you initially add an ICAP scan server to a cluster, the server is automatically enabled, as indicated by a green icon in the **Status** column. If the cluster cannot communicate with the ICAP scan server, the status icon turns red. A disabled ICAP server is indicated by a gray status icon.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Summary**.
The Anti-Virus **Summary** page appears.
2. In the **Actions** column under **ICAP Servers**, click **Enable** for the scanning server that you want to activate on the cluster.
The ICAP scan server's status icon in the **Status** column changes from gray to green.

Disable an ICAP scan server on the cluster

You can disable an ICAP scan server on a cluster if you need to take it offline for maintenance or other reasons.

If other scan servers are enabled on the cluster, disabling one server causes an automatic failover to the remaining scan servers. However, if only one scan server is enabled on the cluster, disabling that server halts any currently running anti-virus scans.



Note: If only one ICAP scan server is configured on the cluster, it is advisable to first disable the cluster anti-virus scanning service before disabling the scan server to prevent unnecessary alerts about unavailable scan servers.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Summary**.
The Anti-Virus **Summary** page appears.

2. In the **Actions** column under **ICAP Servers**, click **Disable** for the scan server that you want to disable on the cluster. The status indicator for the scan server in the **Status** column changes from green to gray.

Test an ICAP server's connection to the cluster

You can test whether an ICAP scan server is reachable from the cluster.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Summary**. The Anti-Virus **Summary** page appears.
2. In the **Action** column under **ICAP Servers**, click **Test connection** for the ICAP scan server that you want to communicate with from the cluster.

Depending on the status of the ICAP scan server that you test, its icon in the **Status** column displays one of the following colors:

- Green: The ICAP server is enabled and responding to the test.
- Green over gray: The ICAP server is disabled and responding to the test.
- Red: The ICAP server is not responding to the test.
- Red over gray: The ICAP server is disabled and not responding to the test.

Pausing the mouse pointer over an icon displays a text message about the ICAP scan server's response to the test.

Modify ICAP scan server settings

If you change the IP address of an ICAP scan server, you must modify that setting on the cluster to ensure that anti-virus scanning can continue.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Summary**. The Anti-Virus **Summary** page appears.
2. In the **Action** column under **ICAP Servers**, click **Edit** for the ICAP scan server whose settings you want to modify. The **Edit ICAP Server** dialog box appears.
3. In the **ICAP URL** box, type the ICAP protocol prefix followed by the IP address of the ICAP scan server. For example: `icap://10.20.100.3`
Optionally, you can append a port number to the URL. Isilon clusters use port 1344 by default for ICAP communications. For example: `icap://10.20.100.3:1344`
Optionally, you can append a path to the URL. For example: `icap://10.20.100.3/avscans`
4. Optionally, in the **Description** box, modify the description for the ICAP scan server. The description can include up to 50 characters.
5. Click **Submit**.

Delete an ICAP server from the cluster

If you delete an ICAP scan server from the cluster, the system halts any currently running anti-virus scans on the cluster if no other ICAP scan servers are configured and enabled.

If other ICAP servers are enabled, any scanning being performed by the deleted ICAP server will fail over to the remaining ICAP servers. When only one ICAP scan server is configured, it is advisable to first disable the cluster anti-virus scanning service before deleting the scan server to prevent unnecessary alerts about unavailable scan servers.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Summary**. The Anti-Virus **Summary** page appears.
2. In the **Actions** column under **ICAP Servers**, click **Delete** for the ICAP scan server that you want to remove from the cluster. A confirmation dialog box appears, asking if you want to delete the ICAP server.

3. Click **Yes**.

Anti-virus global settings

Global settings enable you to specify how all anti-virus scans are performed on the cluster. Some global settings can be overridden by individual scanning policies.

One particularly important anti-virus global setting governs the cluster's response when infected files are detected. third-party ICAP scan servers may be able to repair infected files. If repair is not possible, infected files can be quarantined to prevent end-user access, or truncated to render the threats harmless.


Other global settings enable you to restrict anti-virus scanning to files of up to a specified maximum size, or restrict scans to only files with specific file extensions or specific file names. These settings can be overridden by individual anti-virus scanning policies.

On-access scans govern whether files are scanned for viruses at the time that end users open or close them.

Configure global anti-virus settings

Global settings enable you to specify how all anti-virus scans are performed on the cluster unless those settings are overridden by individual scanning policies.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Settings**.
The Anti-Virus **Settings** page appears.
2. Under **All Scans**, configure the following options that control all anti-virus scans:

Option	Description
Action on Detection	<p>Specifies which action the cluster and ICAP scan server will take if a virus is detected in files:</p> <ul style="list-style-type: none"> • Alert only: Generates an alert at the Warning level when a virus is detected, but does not quarantine or truncate the infected files. • Repair or quarantine: Attempts to repair infected files by sending them to the third-party ICAP scan servers. If repair is not possible, the infected files are quarantined on the cluster so that users cannot access them. • Repair or truncate: Attempts to repair infected files by sending them to the third-party ICAP scan servers. If repair is not possible, the infected files are truncated on the cluster to render them harmless. • Repair only: Attempts to repair infected files by sending them to the third-party ICAP scan servers. If repair is not possible, the cluster generates an alert at the Warning level. • Quarantine: Prevents users from opening or editing infected files. Storage administrators can remove infected files from quarantine through the cluster's File System Explorer. • Truncate: Reduces the size of infected files to zero bytes to render them harmless. File truncation cannot be reversed.
File size restriction	<p>Specifies whether file size is used to determine which files are included in anti-virus scans:</p> <ul style="list-style-type: none"> • Scan all files regardless of size: Includes all files in anti-virus scans regardless of how large they are. • Only scan files smaller than the maximum file size: Excludes files that are larger than a maximum size specified in bytes, megabytes, gigabytes, petabytes, or terabytes. The default setting is to scan files smaller than 2 GB in size. <p> Caution: As a practical limitation, none of the third-party ICAP scan servers currently support scanning of files greater than 2 GB in size. Setting this value to greater than 2GB can cause scans of large files to time out or fail. This can degrade system performance, prevent access to files, or cause policy-based scanning to report spurious failures. You should scan files greater than 2 GB only if you are certain your ICAP servers can successfully scan files this large.</p>

Option	Description
Filename restrictions	<p>Specifies whether to include or exclude files from anti-virus scans based on their file names or extensions:</p> <ul style="list-style-type: none"> • Scan all files: By default, all files are scanned regardless of their names or extensions. • Only scan files with the following extensions or filenames: Restricts scanning to only those files matching the specified file name or extension criteria. • Scan all files except those with the following extensions or filenames: Excludes scanning of files matching the specified file name or extension criteria.
Extensions	If Filename restrictions are enabled, this list identifies which file extensions are either included or excluded from anti-virus scans. Click Edit list to add or modify file extensions, and optionally select from more than 140 commonly used file extensions.
Filenames	If Filename restrictions are enabled, this list identifies which files are either included or excluded from anti-virus scans. Click Edit list to add to or modify the list of files.

3. Under **On Access Scans**, configure the following options to determine whether files are scanned for viruses when they are opened or closed by users:

Option	Description
Scan files when they are opened	Specifies that all files are scanned for viruses when users open them. If files cannot be scanned on open, you can select whether to Allow access or Deny access to users of those files. Scanning files when they are opened causes a delay between the time when users execute a file-open command and when the files are displayed on their computers.
Scan files when they are closed	Scans all open files for viruses when users close those files.
Directories to be scanned	Specifies which directories on the cluster are scanned for viruses when users open or close files. Click Edit list to add to or modify the list of directories.



Note: If Windows file sharing oplocks are enabled, do not enable the **Scan files when they are opened** option. Enabling this option could negatively affect cluster performance. If oplocks are enabled and you select **Scan files when they are opened**, the web administration interface prompts you to disable oplocks. Disabling oplocks will disconnect or disrupt any current SMB client connections.

4. Under **Reports**, select the **Report Retention** option to specify how long in days, weeks, months, or years you want to retain anti-virus scanning reports on the cluster before they are automatically purged.



Note: Some units of time specified for report retention are displayed differently than in the web administration interface than how you originally enter them. Entering a number of days that is equal to a corresponding value in weeks, months, or years results in the larger unit of time being displayed. For example, if you enter a value of **30 days**, the web interface displays that value as "1 month," while **364 days** is represented as "52 weeks," and **365 days** is represented as "1 year." This change occurs because OneFS internally records report retention times in seconds and then converts them into days, weeks, months, or years for display.

5. Click **Submit**.

Anti-virus scanning policies

Anti-virus scanning is controlled by policies that specify which files on the cluster will be scanned and when they will be scanned.

Anti-virus scanning policies can be configured to:

- Scan files in specific root directories on the cluster.

- Run scans at scheduled times on a daily, weekly, monthly, or yearly basis.
- Enable storage administrators to run scans manually at any time.
- Enforce or ignore global anti-virus settings that restrict scans to certain file names, extensions, and maximum file sizes.

Add an anti-virus scanning policy

You can add policies that initiate anti-virus scans on the cluster at scheduled intervals or can be run manually at any time. Policies can either apply or ignore global anti-virus scanning settings.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Policies**.
The Anti-Virus **Policies** page appears.
2. Click **Add Policy**.
The Anti-Virus **Add Policy** page appears.
3. In the **Name** box, type a descriptive name for the anti-virus scanning policy.
The policy name cannot include spaces; however, hyphens and underscore characters are permitted.
4. Next to the **Root directories** box, click **Edit list** to select which root directories on the cluster will be included in the policy's anti-virus scan.
You must select at least one root directory in order to create a policy.
The **Browse OneFS Filesystem** dialog box appears.
5. Click the directories on the cluster that will be included in the policy, and then click **Done**.
You must select at least one root directory in order to create a policy.
6. In the **Restrictions** area, specify whether to enforce global file size and file name restrictions that are configured on the Anti-Virus **Settings** page:
 - Click **Enforce file size and filename restrictions** if you want the policy to apply global restrictions for file names, extensions, and file sizes. You can view these global settings by clicking the **Settings** link to the right of this option.
 - Click **Scan all files within the root directories** if you want to ignore the global restrictions and configure the policy to scan all files in the root directories you selected in steps 4 and 5.
7. For **Run policy**, specify whether the policy's anti-virus scanning will be run **Manually** or **Scheduled** to run automatically at specified intervals.
8. If you selected a **Scheduled** anti-virus scanning policy, configure the following options:
 - a. Under **Interval**, specify whether the policy will run an anti-virus scan on a daily, weekly, monthly, or yearly basis, and then set the appropriate interval period.
 - a. Under **Frequency**, specify whether the policy will run an anti-virus scan **Once** during the specified **Interval**, or **Multiple times**.
9. Click **Submit**.
The new policy appears on the Anti-Virus **Policies** page.

Enable an anti-virus scanning policy

If policies are disabled for any reason, they must be re-enabled before they can be run manually or on a schedule.

When anti-virus scanning policies are created, they are automatically enabled by default. This means that newly added scheduled scans will run at their specified intervals, and that manual scans will run whenever they are triggered by the storage administrator. However, if scheduled policies are disabled for any reason (such as for system maintenance, or to prevent specific anti-virus scans from running), they must be re-enabled before they can be run on a schedule. Similarly, disabled manual scanning policies must be re-enabled before they can be run.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Policies**.
The Anti-Virus **Policies** page appears. Any disabled policies are indicated by gray status icons in the **Run** column.
2. In the **Actions** column under **Policies**, click **Enable** for the policy that you want to enable.

The gray status indicator icon in the **Run** column disappears.

Disable an anti-virus scanning policy

You can disable anti-virus scanning policies to prevent specific scans from running on the cluster according to their schedules. Disabled policies can still be run manually.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Policies**.
The Anti-Virus **Policies** page appears.
2. In the **Actions** column under **Policies**, click **Disable** for the policy that you want to prevent from running according to its schedule.
A gray status icon appears in the **Run** column, indicating that the policy is disabled.

Run a scheduled anti-virus scanning policy manually

When scheduled anti-virus scanning policies are configured and enabled, they will automatically run on the cluster at their scheduled times without requiring storage-administrator intervention.

You can force scheduled policy scans to run manually at any time regardless of the schedule settings. Forcing a scheduled policy scan to run manually does not affect subsequent scheduled scans.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Policies**.
The Anti-Virus **Policies** page appears.
2. In the **Actions** column under **Policies**, click **Start** for the scheduled policy that you want to run manually.
An animated status icon in the **Run** column indicates that the policy anti-virus scan is running.

Run a manual anti-virus scanning policy


Manual anti-virus scanning policies provide the flexibility to perform real-time scans on the cluster.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Policies**.
The Anti-Virus **Policies** page appears.
2. In the **Actions** column under **Policies**, click **Start** for the manual policy that you want to run.
An animated status icon in the **Run** column indicates that the anti-virus scan is running. When the scan completes, a green icon appears in the **Data** column.

Scan for viruses when a file is opened or closed

You can configure the cluster to scan files for viruses when they are opened or closed by end users.

As a safeguard in case an anti-virus scanning server is not available, you can specify whether to permit or deny end-user access to files. When files are scanned at the time end users open them, users will experience a delay in accessing the files while anti-virus scanning is performed.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Settings**.
The Anti-Virus **Settings** page appears.
 2. Under **On Access Scans**, Select one or both of the following check boxes:
 - **Scan files when they are opened:** Performs a virus scan when an end user initiates the opening of a file. If a file cannot be scanned, you can specify whether to **Allow access** or **Deny access** to the file.
 - **Scan files when they are closed:** Performs a virus scan when an end user closes a file.
-  **Note:** Scanning files when users open them should not be enabled when Windows file-sharing oplocks are enabled because of cluster performance considerations. If oplocks are enabled when you select **Scan files when they are opened**, the web administration interface will prompt you to disable oplocks.

3. Specify which directories will be scanned when users open or close files by clicking **Edit list** for **Directories to be scanned**.
The **Browse OneFS Filesystem** dialog box appears.

4. Select the directories on the cluster that will be included in the on-access scans, and then click **Done**.
5. Click **Submit**.

Stop a running virus scan policy

You can stop a currently running anti-virus scanning policy at any time, whether the policy is scheduled or running manually.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Summary**.
The Anti-Virus **Summary** page appears.
2. In the **Action** column under **Currently Running**, click **Cancel** for the running anti-virus scan policy that you want to halt.
In the **Status** field under **Recently Completed**, the halted scan displays an orange icon. Pausing the mouse pointer over the icon displays this message: Scan stopped before completion.
3. Optionally, you can re-start the stopped scan by clicking **Start** for the scan in the **Actions** column.

Modify an anti-virus scanning policy

You can modify configured anti-virus scanning policies as your storage security needs evolve. For example, you could convert a manual anti-virus scan into to a scheduled scan, modify the root directories the scan runs on, or alter a policy scan's schedule.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Policies**.
The Anti-Virus **Policies** page appears.
2. In the **Policy** column, click the name of the anti-virus scanning policy you want to modify.
The Anti-Virus **Edit Policy** page appears.
3. In the **Name** box, type a descriptive name for the anti-virus scanning policy.
Note that the policy name cannot include spaces; hyphens, and underscore characters are permitted.
4. Next to **Root directories**, click **Edit list** to select which root directories on the cluster will be included in the policy's anti-virus scan.
The **Browse OneFS Filesystem** dialog box appears.
5. In the **Browse OneFS Filesystem** dialog box, select which directories on the cluster will be included in the policy, or excluded, and then click **Done**.
6. Use the **Restrictions** options to select whether to enforce global file size and file name restrictions that are configured on the Anti-Virus **Settings** page:
 - Select **Enforce file size and filename restrictions** if you want to apply the global restrictions on which file names, extensions, and file sizes to a policy. You can view these global settings by clicking the **Settings** page link to the right of this option.
 - Select **Scan all files within the root directories** if you want to ignore the global restrictions and have the policy to scan all files in the root directories you selected in steps 4 and 5.
7. For Run policy, select whether the policy's anti-virus scanning will be run **Manually** or **Scheduled** to run automatically at specified intervals.
8. If you specify a **Scheduled** policy of anti-virus scanning, configure the following options:
 - a. Under **Interval**, select whether the policy will run an anti-virus scan on a daily, weekly, monthly, or yearly basis, and then set the appropriate interval period.
 - a. Under **Frequency**, select whether the policy will run an anti-virus scan **Once** during the specified Interval, or **Multiple times**.
9. Click **Submit**.
The modified policy appears on the Anti-Virus **Policies** page.

Delete an anti-virus scan policy

You can delete anti-virus scan policies if you no longer need them. If a policy is currently running an anti-virus scan on the cluster, it cannot be deleted until the scan either completes or is manually stopped.



Note: Deleting an anti-virus scan policy will result in the name of the policy being removed from any existing reports about the policy and replaced with an alphanumeric ID number. To preserve the policy name in anti-virus reports, disable the policy instead of deleting it.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Policies**.
The Anti-Virus **Policies** page appears.
2. In the **Actions** column, click **Delete** for the anti-virus scanning policy that you want to delete.
A dialog box appears and prompts you to confirm whether you want to delete the policy.
3. Click **Yes**.

Anti-virus threat responses

If infected files are detected on the cluster, you can configure how the anti-virus service responds to the threats.

The anti-virus service can respond to threats detected on the cluster in three different ways:

- The third-party ICAP scan servers connected to the cluster may be able to repair some infected files.
- If infected files cannot be repaired, they can be quarantined to prevent end users from accessing them. Quarantined files can be released from quarantine, either to examine them with an external anti-virus tool or if they are found not to be infected.
- Infected files can also be truncated, which reduces the file size to zero bytes and renders the threat harmless.

View detected virus threats on the cluster

You can view all virus threats detected on the cluster on the Anti-Virus **Detected Threats** page. You can also view the most recent threats on the Anti-Virus **Summary** page.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Detected Threats**.
The Anti-Virus **Detected Threats** page appears.
2. The **Detected Threats** page displays the following information about any viruses detected on the cluster:

Option	Description
Status	Displays an icon that indicates the status of the detected threat. Red means a threat was detected in a file but no action was taken. Orange means an infected file was truncated to zero bytes in size. Yellow means an infected file was quarantined.
Threat	Displays the name of the detected threat as it is recognized by the third-party ICAP scan server.
Filename	Displays the name of the infected file.
Directory	Displays the cluster directory in which the infected file is located.
Remediation	Indicates what, if any, remediation was taken against the infected file. Files can be truncated or quarantined based on settings that you configure. If no action was taken against an infected file, the remediation status is shown as infected.
Detected	Displays the date and time when the file was detected.
Policy	Displays the name of the anti-virus scanning policy that detected the threat.
Currently	Displays the current state of the infected file: either truncated, infected, or possibly infected.
File size	Displays the size, in bytes, of the infected file. Files that have been truncated are zero bytes.

Option	Description
Actions	Lists the actions that can be taken against the infected file. Most files can be rescanned, quarantined, or truncated. No actions can be taken against files that have been truncated. You can also view reports for the anti-virus scanning policies that detected the threats.

- Optionally, you can filter the display of detected threats using the **Search** list in the upper right of the page.

By default, the system displays **All threats**, but you can narrow the scope to show only truncated files, quarantined files, or repaired files.

Repair infected files

Infected files detected on the cluster can be repaired automatically by third-party ICAP scan servers.

You can select from among five global file repair options that can be applied to all anti-virus scans whether they are performed manually, run on a schedule, or performed when users open or close files. Repairing an infected file sends the file to the ICAP scan server that detected it, where the repair is performed if possible, and then returned to the cluster. If the ICAP scan server cannot repair an infected file, the file can be quarantined or truncated, or an alert can be generated, depending on your configured settings.

- On the **File System** menu, point to **Anti-Virus**, and then click **Settings**.
The Anti-Virus **Settings** page appears.
- Under **All Scans**, select one of the following **Action on detection** options:
 - Repair or quarantine:** Attempts to repair infected files by sending them to the third-party ICAP scan servers. If repair is not possible, the infected files are quarantined on the cluster so that users cannot access them.
 - Repair or truncate:** Attempts to repair infected files by sending them to the third-party ICAP scan servers. If repair is not possible, the infected files are truncated on the cluster to render them harmless.
 - Repair only:** Attempts to repair infected files by sending them to the third-party ICAP scan servers. If repair is not possible, the cluster generates an alert at the Warning level.
 - Quarantine:** Isolates the infected files on the cluster so that users cannot access them; administrators can still access the files.
 - Truncate:** Reduces the infected files on the cluster to zero bytes in size to render them harmless.
- Click **Submit**.
If an infected file is detected, the ICAP scan server will automatically attempt to repair it. If the repair attempt fails, the infected file will be truncated or quarantined, depending on the **Action on detection** option you specified.

Quarantine an infected file

Infected files detected on the cluster can be quarantined manually or automatically to prevent users from accessing them.

This procedure describes how to manually quarantine a file. To automatically quarantine infected files, configure the anti-virus global settings.

- On the **File System** menu, point to **Anti-Virus**, and then click **Detected Threats**.
The Anti-Virus **Detected Threats** page appears, and displays a list of any detected threats.
- In the **Actions** column, for the infected file that you want to quarantine, click **Quarantine**.
The file's status in the **Currently** column changes to **Quarantined**.

Restore a quarantined file

Infected files that have been quarantined can be restored to repair them with a different tool, or to allow end users to access them.

This is useful, for instance, in the event that a virus scan erroneously identifies a file as being infected, or if you want to attempt to repair an infected file using another application. You can restore a quarantined file from the cluster File System Explorer.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Detected Threats**.
The Anti-Virus **Detected Threats** page appears, and displays a list of any quarantined files. Any quarantined files are flagged as **quarantined** in the **Remediation** column and by a yellow icon in the **Status** column.
2. In the **Filename** column, click the name of the quarantined file that you want to restore.
The **File System Explorer** page appears, and displays information about the infected file. The **Status** of the file is shown as **Quarantined**.
3. Click **Restore**.
The File Properties **Restore** page appears.
4. Select the **Restore file from quarantine** check box, and then click **Submit**.
The File System Explorer **File Properties** page appears. Note that quarantine status information about the file is no longer displayed.

Truncate an infected file

Any infected files detected on the cluster can be truncated manually or automatically, which renders infections harmless by reducing files to zero bytes.

This procedure describes how to manually truncate an infected file. To automatically truncate infected files, configure the anti-virus global settings.



Note: After an infected file has been truncated, it cannot be repaired or restored to its original size, except through snapshots or by restoring the file from a backup.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Detected Threats**.
The Anti-Virus **Detected Threats** page appears, and displays a list of infected files.
2. In the **Actions** column, click **Truncate** for the file that you want to truncate.
A confirmation dialog box appears and prompts you to confirm whether you want to truncate the infected file.
3. Click **Yes**.

Anti-virus scan reports

You can view anti-virus reports that contain summary and detail information about anti-virus scans run on the cluster.

You can also export anti-virus scan reports as comma-separated values (.csv) files. Any virus threats detected on the cluster are also reported as alerts, as are problems with the availability of third-party ICAP scan servers.

You can configure global anti-virus settings to specify how long to retain anti-virus scanning reports on the cluster before they are automatically purged.

View anti-virus scan reports

You can view reports about anti-virus scans performed on the cluster. Report information includes the start time and duration of anti-virus scans, number of files scanned, number of threats detected, and details about the threats.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Reports**.
The Anti-Virus **Reports** page appears, and displays summary information about anti-virus scans performed on the cluster.
2. Optionally, filter the report information using the two **Search** lists in the upper right of the page:
 - The left list controls which types of anti-virus policies the report covers. The default setting is **All policies**, but you can narrow the report scope to show scan results for a specific policy, for manual scans, for scans on opened files, or for scans on closed files.
 - The right list controls the time period the displayed report covers. The default setting is **Last week**, but you can narrow the report scope to time periods ranging from the last day of anti-virus scans to all scans run during the past two years, or select **Any time** to show all historical report data.
3. To view more information about a specific scan's results, under **Actions** click **View details**.

The Anti-Virus **Report Detail** page appears. It provides details about specific threats detected by the scan, as well as the scan's start and end times, scan duration in hours:minutes:seconds, number of files scanned, and bandwidth of traffic sent from the cluster to the ICAP scan server measured in bits per second.

Export an anti-virus report

You can export reports for individual anti-virus scans as comma-separated values (.csv) files. This can be useful for analyzing anti-virus scan data using an external application.

1. On the **File System** menu, point to **Anti-Virus**, and then click **Reports**.
The Anti-Virus **Reports** page displays.
2. In the **Actions** column, click **Export** for the anti-virus policy whose scan results you want to export.
A dialog box prompts you to save or open the .csv file.
3. Specify whether to save or open the .csv file, and then click **OK**.

View anti-virus alerts

If an anti-virus scan detects a threat on the cluster, it triggers a warning-level alert. If no third-party ICAP scan servers are configured or communicating with the cluster, this will also trigger an alert.

1. On the **Status** menu, click **Alerts**.
The **Alerts** page appears.
2. Review the **Active** and **Historical** alerts sections for information on detected threats:
 - Alerts for detected virus threats are classified as warnings and contain the phrase `Anti-Virus scan found threats`. These alerts do not provide threat details, but instead refer to specific reports on the Anti-Virus **Reports** page.
 - Alerts for unavailable ICAP scan servers are classified as warnings and contain the phrase `No ICAP Servers available`. You can view, configure, and test the connections to ICAP scan servers on the Anti-Virus **Summary** page.
 - Alerts for unresponsive ICAP scan servers are classified as warnings and contain the phrase `ICAP Server Unresponsive or Invalid`, and provide the IP address of the server. You can view, configure, and test the connections to ICAP scan servers on the Anti-Virus **Summary** page.

Authentication, identity management, and authorization

OneFS supports a mixed environment in which both Windows Access Control Lists (ACLs) and standard UNIX permissions can be configured on the cluster file system. OneFS provides advanced identity management options to enable proper access controls.

Windows and UNIX permissions cannot coexist on a single file or directory; however, the Isilon file system can translate between Windows and UNIX permissions on the fly.

Isilon supports the following authentication methods:

- Anonymous (requires no authentication)
- Local (queries a database of users and groups stored on the Isilon cluster)
- File (uses a standard UNIX `/etc/passwd` file)
- External authentication:
 - Microsoft Active Directory
 - Lightweight Directory Access Protocol (LDAP)
 - Network Information Service (NIS)

Authentication sources

OneFS supports a variety of authentication sources to verify a user's credentials. Based on the results of the authentication process, the system allows or denies access to view or modify stored data through the various file sharing protocols it supports.

The following authentication sources are supported:

- Local database
- File databases
- Microsoft Active Directory
- Lightweight Directory Access Protocol (LDAP)
- Network Information Service (NIS)



Note: OneFS supports the use of more than one concurrent authentication source. It is important that you fully understand their interactions before enabling multiple authentication sources on the cluster.

View authentication service status

The Authentication Sources **Summary** page displays the status and configured settings for the Active Directory, LDAP, and NIS services.

On the **File Sharing** menu, point to **Authentication Sources**, and then click **Summary**.

The Authentication Sources **Summary** page appears, and displays the following sections:

- **Active Directory:** Displays the Active Directory configuration status and local workgroup or domain settings.
- **LDAP:** Displays the service state and configured settings for the LDAP service.
- **Legacy LDAP:** Displays the service state and configured settings for the legacy version of the LDAP service.

- **NIS:** Displays the service state and configured settings for the NIS service.

Local users and groups

You can configure user and group accounts to control access to files and directories that are stored on the Isilon cluster.

Local authentication can be useful when Active Directory, LDAP, or NIS directory services are not used, or when a specific user or application needs to access the cluster.

View users

You can view the username, full name, group memberships, home directory, and shell settings for local and remote users.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **Local Users**. The Authentication Sources **Local Users** page appears.
2. Optionally click the **Show remote users** link to display system, NIS, and LDAP users.



Note: Remote users may take several minutes to display, depending on the number of users in the NIS and LDAP databases. To hide remote users, click the **Hide remote users** link.

3. Review the following user settings:

- **User:** Displays the username of the user account.
- **Full Name:** Displays user's full name.
- **Group:** Displays the primary group to which the user is assigned.
- **Additional Groups:** Displays any additional groups to which the user is assigned.
- **Home:** Displays the user's home directory.
- **Shell:** Displays the user's shell.
- **Remote:** Displays **Yes** if the user is from a remote source (NIS or LDAP) and cannot be modified on the cluster. By default, only local users are displayed.



Note: You can click a username to view additional settings or to modify a local user's settings.

View groups

You can view user membership for local and remote groups.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **Local Groups**. The Authentication Sources **Local Groups** page appears.
2. Review the following group settings:
 - **Group:** Displays the name of the group.
 - **Members:** Displays any users that are assigned to the group.
 - **Remote:** Displays **Yes** if the group is from a remote source and cannot be modified on the cluster. By default, only local groups are displayed.



Note: You can click a group name to view additional settings or to modify the members of a local group.

Create a local user

When creating a local user account, you can configure its name, password, home directory, UID, shell, and group membership settings. You can optionally disable the account or prevent the password from expiring.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **Local Users**. The Authentication Sources **Local Users** page appears.

2. Click the **Add user** link.
The **Add Local User** page appears.
3. Configure the following settings:
 - **User name:** Type a username for the account.
 - **Full name:** Type a full name for the user.
 - **Password:** Type a password for the account.
 - **Confirm password:** Re-type the password for the account.
 - **Home directory:** Optionally type the full path to the user's home directory. If you do not specify a path, a directory is automatically created at `/ifs/home/<Username>`.
 - **User ID:** Optionally type a UNIX user ID (UID). The default setting is **(auto)**, which assigns the next available UID.
 - **Primary group:** Optionally click to select a primary group from the menu. If you select **(auto)**, a new group is created with the same name as the username. By default, the primary group is **Isilon Users** (GID 1800).
 - **Shell:** This setting applies only to users who will access the file system through SSH or Telnet. Optionally click to select a shell from the list (**sh**, **csch**, **tcsh**, **bash**, **rbash**, **zsh**, or **nologin**). By default, the **zsh** shell is selected.
 - **Enabled:** Optionally select this check box to allow the user to authenticate against the local database for SSH, FTP, Telnet, and Windows file sharing via Server Message Block (SMB). This setting is not used for HTTP file sharing or UNIX file sharing via network file system (NFS).
 - **Password expiry:** Optionally select this check box to prevent the password from expiring.
4. Optionally assign the user to one or more additional groups.
 - a. For the **Additional groups** setting, click **Add**.
The **Choose group(s)** dialog box appears.
 - b. Optionally type a value in the **Name** or **Description** text box to filter the search results.
 - c. Click **Search**.
 - d. Click to select one or more groups from the **Search Results** list, and then click **Choose**.
The dialog box closes, and the **Additional Groups** list is updated.



Note: To remove the user from a group, select the check box next to the group name and then click **Delete**; to remove the user from multiple groups, select the check box for each group and then click **Delete selected**.

5. Click **Submit**.

Modify a local user

You can modify any setting for an existing local user account except its user name.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **Local Users**.
The Authentication Sources **Local Users** page appears.
2. In the list of existing users, click the name of the local user whose settings you want to modify.



Note: Although you can choose to display remote users, only local users can be modified.

The **Modify Local User** page appears.

3. Modify the user's settings as needed.



Note:

- To keep the existing password, leave the **Password** and **Confirm password** boxes empty.
- If you modify the **Home directory** setting, the existing directory remains, and its contents are not copied to the new directory.

4. Click **Submit**.

Delete a local user

When you delete a local user account, the existing home directory remains untouched.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **Local Users**.
The Authentication Sources **Local Users** page appears.
2. In the list of existing users, click the name of the local user that you want to delete.



Note: You can only delete local users.

The **Modify Local User** page appears.

3. Click **Delete User**.

Create a local group

You can create local groups and assign local users to them through the user interface.



Note: Only local users can be added to a group through the web-based user interface. You can modify groups or well-known SIDs by running the `isi auth local groups modify` command.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **Local Groups**.
The Authentication Sources **Local Groups** page appears.
2. Click the **Add group** link.
The **Add Local Group** page appears.
3. Configure the following settings:
 - **Group name:** Type a name for the group.
 - **Group ID:** Optionally type a UNIX group ID (UID). The default setting is (**auto**), which assigns the next available GID.
 - **Users:** Optionally click the check box for each local user that you want to add to the group.
4. Click **Submit**.

Modify a local group

You can add or remove users from local groups.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **Local Groups**.
The Authentication Sources **Local Groups** page appears.
2. In the list of existing groups, click the name of the group that you want to modify.



Note: You can only modify local groups.

The **Modify Local Group** page appears.

3. To change the GID, type a new value in the **Group ID** box. You must choose a unique value.
4. In the **Users** list, select or clear users as needed.
5. Click **Submit**.

Delete a local group

Perform the following steps to delete a local group.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **Local Groups**.
The Authentication Sources **Local Groups** page appears.
2. In the list of existing groups, click the name of the local group that you want to delete.



Note: You can only delete local groups.

The **Modify Local Group** page appears.

3. Click **Delete Group**.

File provider

OneFS uses `/etc/spwd.db` and `/etc/group` files for identities associated with running and administering the cluster. These files do not include information about the identities of end users; however, you can use the file provider to manage end-user identity information based on the format of these files.

The file provider enables you to provide an authoritative third-party source of user and group information to the cluster.

The file provider uses two file formats employed by the BSD operating system.

- The `spwd.db` format provides fast access to the data in the `/etc/master.passwd` file.
- The `/etc/group` format is available in most UNIX operating systems.

Configure the file provider

The file provider pulls directly from two files formatted in the same manner as `/etc/group` and `/etc/passwd`. Updates to the files can be scripted.



Note: To ensure that all nodes in the cluster have access to the same version of the file provider files, it is recommended that you put them within the `/ifs/.ifsvar` directory.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **File Provider**.
The Authentication Sources **File Provider** page appears.
2. Modify the file provider settings as needed.



Note: To activate a setting, you must select its check box. You can deactivate a setting without removing it by clearing its check box.

- **Group file:** Type the full path to the group replacement file.
- **Passwd file:** Type the full path to the supplemental password file.

3. Click **Submit**.

File provider password database format

The file provider password database uses the `spwd.db` format for fast access to the data in the `/etc/master.passwd` file.

The source password file has ten colon-separated fields. A sample line looks like the following:

```
www:*:80:80::0:0:World Wide Web Owner:/nonexistent:/usr/sbin/nologin
```

The fields are defined below in the order in which they appear in the file:

- **Username:** The user's name. This field is case sensitive. OneFS does not set a limit on the length; however, many applications truncate the name at 16 characters.
- **Password:** The user's password in crypt format. If authentication is not required for the user, an asterisk (*) can be substituted in place of a password. The * character is guaranteed to not match any password.
- **UID:** The user's primary identifier. This value should be in the range of **0-4294967294**. Take care when choosing a UID to ensure that it does not conflict with an existing account. For example, do not choose the reserved value **0** as the UID.



Note: There is no guarantee of compatibility if an assigned value conflicts with an existing UID.

- **GID:** The group identifier of the user's primary group. All users are a member of at least one group, which is used for access checks and can also be used when creating files.
- **Class:** This field is not supported by OneFS and should be blank.
- **Change:** Password change time. OneFS does not support changing passwords of users in the file provider.
- **Expire:** Expiration time. OneFS does not support expiration time of users in the file provider.
- **Gecos:** Gecos information. This field can store a variety of information, although it is usually used to store the user's full name.
- **Home:** The user's home directory. This field should point to a directory on `/ifs`.
- **Shell:** The absolute path to the user's shell (**sh**, **csh**, **tcsh**, **bash**, **rbash**, **zsh**, or **nologin**). For example, to deny command-line access to the user, set the shell to `/usr/sbin/nologin`.

UNIX systems often define the `passwd` format as a subset of the above fields, omitting the class, change, and expire fields. To convert a file from `passwd` format to `master.passwd`, add `:0:0:` between the GID field and the Gecos field.

File provider group database format

The file provider uses the format of the `/etc/group` file found in most UNIX operating systems.

The group file is composed of four colon-separated fields. A sample line looks like the following:

```
admin:*:10:root,admin
```

The fields are defined below in the order in which they appear in the file:

- **Group name:** The group's name. This field is case sensitive. Although OneFS does not set a limit on the length, many applications truncate the name at 16 characters.
- **Password:** This field is not supported by OneFS and should be set as an asterisk (*).
- **GID:** The group identifier. This value should be in the range of **0-4294967294**. Be careful when choosing a GID to ensure that it does not conflict with an existing group.
- **Group members:** A comma-delimited list of user names that make up the group's members.

Active Directory

Active Directory directory service is a Microsoft implementation of Lightweight Directory Access Protocol (LDAP), Kerberos, and DNS technologies that can store information about network resources. Active Directory can serve many functions, but the primary reason for joining the cluster to an Active Directory domain is to perform user and group authentication.

When the cluster joins an Active Directory domain, a single Active Directory machine account is created. The machine account is used to establish a trust relationship with the domain and to enable the cluster to authenticate and authorize users in the Active Directory forest. By default, the machine account is named the same as the cluster; however, if the cluster name is more than 15 characters long the name is hashed and is displayed after joining the domain.

View Active Directory settings

The Authentication Sources **Active Directory** page allows you to view and configure local and domain server mode settings. If Active Directory is configured, the current connection status between the cluster and the Active Directory domain is displayed.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **Active Directory**. The Authentication Sources **Active Directory** page appears, and displays the following sections:
 - **Server Status:** Displays whether the cluster is currently joined to an Active Directory domain.
 - A **green** icon indicates that Active Directory is enabled and configured.
 - A **red** icon indicates that Active Directory is enabled and offline.
 - A **gray** icon indicates the server is in local mode (Active Directory is not set up).

If the cluster is joined to an Active Directory domain, the configuration settings are displayed.

- **Server Mode:** Displays the current server mode and allows you to view and configure server mode settings.
- **Cluster name:** Displays the name that the cluster presents to the Active Directory server. When the cluster joins the Active Directory domain, a single machine account of the same name is created on the domain controller.
- **Server description:** Specifies a description of the Isilon cluster to provide to the Active Directory server. The server description is visible to Active Directory server administrators.
- **Server mode:** Specifies the mode of the Isilon server's relationship with Active Directory.

2. Optionally click the **Show advanced settings** link to view and manage additional configuration settings.

Configure the Active Directory service

Active Directory can be configured to run in either local mode or domain mode.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **Active Directory**.
The Authentication Sources **Active Directory** page appears.
2. Optionally, in the **Server description** text box, type a description of the Isilon cluster to provide to the Active Directory server. The server description is visible to Active Directory server administrators.
3. For the **Server mode** setting, click to select either **Local Mode** or **Domain Mode**.
 - For local mode, optionally type the name of a workgroup or keep the default WORKGROUP setting.
 - For domain mode, specify the following settings:
 - **DNS domain:** Type the DNS domain name of the Active Directory server to connect to.
 - **Domain account username:** Type the username of an account that is authorized to join the Active Directory server.
 - **Domain account password:** Type the password for the specified domain account username.
 - **Organizational unit:** Optionally specify an organizational unit (OU), or container, to connect to on the Active Directory server. Domain join will fail if the machine account exists but resides in another OU.
4. Optionally click **Show advanced settings** to view and set additional domain mode options.



Note: These options will be saved only after a successful domain join.

5. Click **Submit**.

Kerberos

The OneFS implementation of the Kerberos version 5 (v5) protocol is optimized to use an Active Directory keytab for Kerberos ticket generation, although it is possible to use a third-party keytab. By default, Active Directory is configured as the Kerberos service for both SMB and NFS.

To accept Kerberos tickets, a Kerberos keytab is required. A keytab contains a collection of keys of various types organized by their principal name (for example, name@domain). OneFS treats all keytab entries as equals; the principal in the request is not matched with the keytab entry principal. OneFS only checks to ensure that it can decrypt the ticket and that the ticket is valid.



Note: When a domain password is changed, a custom memory keytab is used to automatically read the new password and keep the old password available for a limited time. If you restart the lwiod or lsassd daemon soon after changing a domain password, it is possible for the cluster to deny clients with the old (but valid) password until they refresh their cache.

Configure Kerberos

Kerberos configuration is supported via the command line only. Most settings require modification only if using a Kerberos Key Distribution Center (KDC) other than Active Directory.

In addition to the global Kerberos configuration file, OneFS includes a Kerberos configuration file for Active Directory. You can modify either file by following this procedure.

1. Run the `isi auth config krb5` command with the `add`, `modify`, or `delete` subcommand to specify which entries to modify in the Kerberos configuration file.

For a complete list of available subcommands and options, see the OneFS 6.5 Command Reference.

2. Propagate the changes to the Kerberos configuration file by running the `isi auth config krb5 write` command.



Note: By default, changes are written to the global Kerberos configuration file, `/etc/krb5.conf`. To update the Kerberos configuration file for Active Directory, use the `--path` option to specify the `/etc/likewise-krb5-ad.conf` file.

LDAP

The Lightweight Directory Access Protocol (LDAP) is a networking protocol that enables you to define, query, and modify directory services and resources.



Note: The LDAP provider has been updated for this release of the OneFS operating system. If you upgraded from an earlier OneFS version with LDAP enabled, your old LDAP settings are retained on the **Legacy LDAP** page. Legacy LDAP support will be removed in a future release of the OneFS operating system.

A key advantage of LDAP is the open nature of the directory services and the ability to use LDAP across many platforms. The Isilon clustered storage system can authenticate users and groups against an LDAP repository in order to grant them access to the cluster.

View LDAP settings

The Authentication Sources **LDAP** page displays the server status and configured settings for the LDAP provider.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **LDAP**.
The Authentication Sources **LDAP** page appears, and displays the following sections:
 - **Server Status:** Displays the current state of the LDAP service (enabled, disabled, or disabled and not configured) and the LDAP server configuration status.
 - **LDAP Provider Settings:** Displays the currently configured LDAP settings.
 - **Attribute map:** Displays the attribute map (**rfc2307**, **ad**, or **ldapsam**), which defines the default attribute settings to use with the LDAP server.
 - **Authentication enabled:** Displays whether the LDAP provider is enabled to respond to authentication requests.
 - **Balance servers:** Displays whether LDAP servers are affinitized randomly (**yes**) or according to the order in which they are listed in the **Server URI(s)** setting (**no**).
 - **Base distinguished name:** Displays the distinguished name of the entry in the directory from which searches initiated by LDAP clients occur. Also referred to as the search base or base DN.
 - **Bind distinguished name:** Displays the distinguished name of the entry used to bind to the LDAP server.
 - **Bind password:** Displays the password that is used when binding to the LDAP server. Use of this password does not require a secure connection; if the connection is not using TLS the password will be sent in clear text.
 - **Bind timeout period:** Displays the timeout period, in seconds, after which binding to the current server will not be retried and will fail.
 - **Server check interval:** Displays the interval, in seconds, at which checks are made for valid and invalid servers.

- **Provider enabled:** Displays whether the LDAP provider is enabled (**yes**) or disabled (**no**).
- **Enumerate groups:** Displays whether the LDAP provider is able to respond to getgrent requests.
- **Enumerate users:** Displays whether the LDAP provider is able to respond to getpwent requests.
- **Ignore TLS errors:** Displays whether the LDAP provider will ignore all transport layer security (TLS) errors. When enabled, this setting allows administrators to use TLS with an expired or self-signed server certificate.
- **Secure connection required:** Displays whether a secure connection is required when binding with a password to authenticate a user or retrieving password-related attributes.
- **Search timeout period:** Displays the timeout period, in seconds, after which a search will not be retried and will fail.
- **Server URI(s):** Displays a comma-separated list of LDAP server URIs. If the **Balance servers** option is set to **no**, servers will be accessed in the order in which they are listed.

2. To view or adjust individual attribute mappings, click **Show advanced settings**. These mappings correspond to the active **Attribute map** setting.

Configure the LDAP service

To enable the LDAP service, you must configure a base distinguished name (base DN) and at least one LDAP server.



Note:

- If you upgraded from a previous version of OneFS with LDAP enabled, your settings are retained on the **Legacy LDAP** page. You should disable the legacy LDAP service before you configure the updated LDAP service.
 - Specifying an LDAP server that is inaccessible can trigger timeouts and other errors when the cluster attempts to contact it.
 - If the LDAP service is currently enabled, removing all servers or the base distinguished name will disable the LDAP service.
1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **LDAP**.
The Authentication Sources **LDAP** page appears.
 2. In the **LDAP Provider Settings** section, choose an **Attribute map** setting. The attribute map defines the default attribute settings to use with the LDAP server. Changing this setting alters the attribute map available under **Show advanced settings**.
 - **rfc2307:** Uses RFC 2307 default attributes.
 - **ad:** Uses Microsoft Services for Unix (SFU) default attributes for Active Directory integration.
 - **ldapsam:** Uses Samba extension default attributes.
 3. Configure the following settings as needed. In order to change a setting, you must first select its check box.
 - **Authentication enabled:** Select **yes** to enable the LDAP provider to respond to authentication requests, or select **no** to prevent responding to authentication requests.
 - **Balance servers:** Select **yes** to affinitize to a random server, or **no** to affinitize according to the order in which the servers are listed in the **Server URI(s)** setting.
 - **Base distinguished name:** Type the distinguished name of the entry at which to start LDAP searches. Base DN's may include **cn** (Common Name), **l** (Locality), **dc** (Domain Component), **ou** (Organizational Unit), or other components. An example base DN might be **dc=isilon,dc=com**.
 - **Bind distinguished name:** Type the distinguished name of the entry to use to bind to the LDAP server.
 - **Bind password:** Specify the password to use when binding to the LDAP server. Use of this password does not require a secure connection; if the connection is not using TLS the password will be sent in clear text.
 - **Bind timeout period:** Specify the number of seconds after which binding to a server will not be retried and will fail. The default value is 60.
 - **Server check interval:** Specify the number of seconds between checking for valid and invalid servers. The default value is 180.

- **Provider enabled:** Select **yes** to enable, or **no** to disable, the LDAP provider.
- **Enumerate groups:** Specify whether to allow the LDAP provider to respond to getgrent requests.
- **Enumerate users:** Specify whether to allow the LDAP provider to respond to getpwent requests.
- **Ignore TLS errors:** Specify whether to ignore all transport layer security (TLS) errors. When enabled, this setting allows administrators to use TLS with an expired or self-signed server certificate.
- **Secure connection required:** Specify whether to require a secure connection when binding with a password or retrieving password-related attributes.
- **Search timeout period:** Specify the number of seconds after which a search will not be retried and will fail. The default value is 100.
- **Server URI(s):** Type one or more valid LDAP server URIs, separated by commas, in the format `ldap://server:port` (LDAP) or `ldaps://server:port` (secure LDAP).

**Note:**

- If you do not specify a port, the default port is used (389 for LDAP; 636 for secure LDAP).
- If non-secure LDAP (`ldap://`) is specified, the Bind password will be transmitted to the server in clear text.
- If the **Balance servers** option is set to **no**, servers will be accessed in the order in which they are listed.

4. Optionally click **Show advanced settings** to fine-tune the LDAP configuration settings. These settings vary according to the selected **Attribute map** setting.
5. Click **Submit**.

The User Authentication **LDAP** page refreshes. If the **Base distinguished name** and **Server URI(s)** settings are configured and **Provider enabled** is set to **yes**, LDAP is enabled automatically.



Note: If the LDAP service and the legacy LDAP service are both enabled, a warning message displays. Follow these steps to disable the legacy LDAP service:

1. Click **Legacy LDAP** at the top of the page.
2. On the **Legacy LDAP** page, under **Server Status**, click **Disable**.

Enable or disable the LDAP service

The LDAP service is automatically enabled when it is configured. You can manually disable or enable the LDAP service by setting the LDAP **Provider enabled** attribute.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **LDAP**.
The Authentication Sources **LDAP** page appears.
2. In the **LDAP Provider Settings** section, change the state as needed.
 - To disable the LDAP service, click the **Provider enabled** check box and select **no**.
 - To enable the LDAP service, click the **Provider enabled** check box and select **yes**.

Legacy LDAP

The OneFS operating system has been updated with a newer version of the LDAP provider. If you have upgraded your cluster from an earlier version of OneFS, any existing LDAP settings are now classified as 'legacy LDAP' and you must manually switch to the newer version. Legacy LDAP support will be removed in a future release.

View legacy LDAP settings

The Authentication Sources **Legacy LDAP** page displays the state and configured settings for the legacy LDAP service.



Note: If the legacy LDAP service and the newer LDAP service are both enabled, a warning message displays. Disable the legacy LDAP service by clicking **Disable** in the **Server Status** section.

On the **File Sharing** menu, point to **Authentication Sources**, and then click **Legacy LDAP**. The Authentication Sources **Legacy LDAP** page appears, and displays the following sections:

- **Server status:** Displays the current state of the legacy LDAP service (enabled, disabled, or disabled and not configured).
- **Legacy LDAP Settings:** Displays the currently configured settings for the legacy LDAP provider.
 - **PAM Authentication:** Displays whether Pluggable Authentication Modules (PAM) is enabled for authentication. PAM authentication is disabled by default.
 - **Base distinguished name:** Also referred to as the search base or base DN, identifies the entry in the directory from which searches initiated by LDAP clients occur.
 - **Port:** Displays the LDAP port (389, by default).
 - **Legacy LDAP servers:** Displays a list of configured LDAP servers in the order they will be accessed.
 - **Legacy LDAP extra settings:** Displays any extra LDAP parameters that are configured.

Configure the legacy LDAP service

Legacy LDAP support will be discontinued in a future release. We recommend using the newer LDAP service, available at the Authentication Sources **LDAP** page of the web administration interface.



Note: You must configure the base distinguished name (base DN), port number, and at least one LDAP server before you can enable the legacy LDAP service.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **Legacy LDAP**. The Authentication Sources **Legacy LDAP** page appears.
2. In the **Legacy LDAP Settings** section, click **Edit**. The **Configure LDAP Settings** dialog box appears.
3. Set the **Base distinguished name** value. The base distinguished name (base DN), commonly referred to as the search base, identifies the entry in the directory from which searches initiated by LDAP clients occur. Base DNs may include `cn` (Common Name), `l` (Locality), `dc` (Domain Component), `ou` (Organizational Unit), or other components. An example base DN might be `dc=isilon,dc=com`.
4. In the **Port** text box, enter **389** (the default port) or an alternative port number that LDAP services will use.
5. Optionally select the **PAM Authentication** checkbox to enable LDAP authentication through Pluggable Authentication Modules (PAM).



Note: By default, the cluster does not use LDAP for authentication; only user and group information is pulled. Leave this option cleared to prevent authentication with the LDAP source, for example in NFS-only configurations or if accounts exist in multiple authentication sources.

6. In the **LDAP servers** list, add, reorder, or remove LDAP servers as needed. LDAP servers are accessed in the order in which they are listed.
 - To add an LDAP server to the list, click **New** and then type the IP address or hostname of the server.
 - To change the order in which an LDAP server is accessed, click its IP address or hostname and then click the **up** or **down** arrow.
 - To remove an LDAP server from the list, click its IP address or hostname and then click **Delete**.



Note:

- At least one valid server IP address or host name must be specified in order to configure and enable the legacy LDAP service.
- Specifying an LDAP server that is inaccessible can trigger timeouts and other errors when the cluster attempts to contact it.
- If the legacy LDAP service is currently enabled, removing all servers will disable the legacy LDAP service.

7. Optionally modify the **Extra LDAP parameters** list.

- To add an LDAP parameter to the list, click **New** and then fill in the **Key** and **Value** text boxes.
- To modify an existing parameter, click its key or value to overwrite its current setting.
- To remove an LDAP parameter from the list, click its key or value and then click **Delete**.

8. Click **Submit**.

The User Authentication **Legacy LDAP** page appears.



Note: If all required settings are configured, and the newer LDAP service is not enabled, you can enable the legacy LDAP service by clicking **Enable** in the **Server Status** section.

Enable or disable the legacy LDAP service

You must configure the legacy LDAP service before you can enable it. Do not enable the legacy LDAP service if the newer LDAP service is enabled.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **Legacy LDAP**.

The Authentication Sources **Legacy LDAP** page appears.

2. In the **Server Status** section, change the state as needed.

- If the legacy LDAP service is disabled and not configured, you must configure the settings before you can enable it.
- If the legacy LDAP service is currently disabled, you can enable it by clicking **Enable**.



Note: Only one LDAP provider should be running at a time. You should not enable the legacy LDAP provider unless the newer LDAP provider is disabled.

- If the legacy LDAP service is currently enabled, you can disable it by clicking **Disable**.

NIS

The Network Information Service (NIS) provides authentication and uniformity across local area networks. The Isilon cluster includes a NIS component that enables you to integrate the cluster into an existing NIS infrastructure in your network.



Note: The NIS provider has been updated for this release of the OneFS operating system. NIS is no longer used for hostname lookups by default. If you upgraded from an earlier OneFS version with NIS enabled and you wish to use NIS for hostname lookups, you must manually configure the hostname resolution order.

NIS, designed by Sun Microsystems, is a directory services protocol that can be used by the Isilon clustered storage system to authenticate users and groups when accessing the cluster. NIS stores information about the network, workstation names and addresses, users, and network services. Originally known as yp (Yellow Pages), NIS is different from NIS+, which the Isilon cluster does not support.

NIS provides generic database access capabilities that enable you to distribute information, such as that contained in the password and groups files, to all hosts on your network. This makes the network appear as if it were a single system, with the same accounts on all hosts.

View NIS settings

The Authentication Sources **NIS** page displays the state and configured settings for the NIS service.

On the **File Sharing** menu, point to **Authentication Sources**, and then click **NIS**.

The Authentication Sources **NIS** page appears, and displays the following sections:

- **Server Status:** Displays the state of the NIS service and NIS server configuration.
- **Host Resolution Settings:** Displays the order in which DNS and NIS host name resolution occurs if NIS is enabled.

- **NIS Provider Settings:** Displays currently configured settings for the NIS provider.
 - **Authentication enabled:** Displays whether the provider is enabled to respond to authentication requests.
 - **Balance servers:** Displays whether NIS servers are affinitized randomly (**yes**) or according to the order in which they are listed in the **Server(s)** setting (**no**).
 - **Server check interval:** Displays the interval, in seconds, at which checks are made for valid and invalid servers.
 - **Domain:** Displays the NIS server domain name, if configured.
 - **Provider enabled:** Displays whether the NIS provider is enabled (**yes**) or disabled (**no**).
 - **Enumerate groups:** Displays whether the NIS provider is able to respond to getgrent requests.
 - **Enumerate users:** Displays whether the NIS provider is able to respond to getpwent requests.
 - **Request timeout period:** Displays the timeout period, in seconds, after which a request will not be retried and will fail.
 - **Request retry period:** Displays the timeout period, in seconds, after which a request will be retried.
 - **Server(s):** Lists any NIS servers that are currently configured, in the order in which they will be used.

Configure the NIS service

To enable the NIS service, you must configure a NIS domain and at least one NIS server.



Note:

- If you upgraded from a previous version of OneFS with NIS enabled, in order to use NIS for hostname lookups you must manually specify the hostname resolution order.
 - At least one NIS server should be accessible by all nodes on the cluster; otherwise, delays and/or timeouts may occur while performing tasks that query information from NIS.
 - If the NIS service is currently enabled, removing all servers or the domain setting will disable the NIS service.
1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **NIS**.
The Authentication Sources **NIS** page appears.
 2. In the **Host Resolution Settings** section, select the hostname resolution order.
 - **NIS disabled:** Does not use NIS for hostname resolution.
 - **DNS before NIS:** Attempts hostname resolution through DNS first, then through NIS. This setting is recommended only if your NIS server contains DNS records.
 - **NIS before DNS:** Attempts hostname resolution through NIS first, then through DNS. This setting is recommended only if your NIS server contains DNS records.
 3. In the **NIS Provider Settings** section, configure the following settings as needed. In order to change a setting, you must first select its check box.
 - **Authentication enabled:** Select **yes** to enable the NIS provider to respond to authentication requests, or select **no** to prevent responding to authentication requests.
 - **Balance servers:** Select **yes** to affinitize to a random server, or **no** to affinitize according to the order in which the servers are listed in the **Server(s)** setting.
 - **Server check interval:** Specify the number of seconds between checking for valid and invalid servers. The default value is 180.
 - **Domain:** Specify the NIS server domain name.
 - **Provider enabled:** Select **yes** to enable, or **no** to disable, the NIS provider.
 - **Enumerate groups:** Specify whether to allow the NIS provider to respond to getgrent requests.
 - **Enumerate users:** Specify whether to allow the NIS provider to respond to getpwent requests.
 - **Request timeout period:** Specify the number of seconds after which a request will not be retried and will fail. The default value is 20.

- **Request retry period:** Displays the timeout period, in seconds, after which a request will be retried. The default value is 5.
- **Server(s):** Type one or more valid IP addresses, hostnames, or fully qualified domain names, separated by commas.



Note: If the **Balance servers** option is set to **no**, servers will be accessed in the order in which they are listed.

4. Click **Submit**.

The User Authentication **NIS** page refreshes. If the **Domain** and **Server(s)** settings are configured and **Provider enabled** is set to **yes**, NIS is enabled automatically.

Enable or disable the NIS service

The NIS service is automatically enabled when it is configured. You can manually disable or enable the NIS service by setting the NIS **Provider enabled** attribute.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **NIS**.
The Authentication Sources **NIS** page appears.
2. In the **NIS Provider Settings** section, change the state as needed.
 - To disable the NIS service, click the **Provider enabled** check box and select **no**.
 - To enable the NIS service, click the **Provider enabled** check box and select **yes**.

Advanced authentication settings

Advanced authentication settings include ID mapping, on-disk identity, and other general settings.



Note: This section describes the settings in the **General Settings** area only.

- For details about configuring the settings in the **ID Mapping Settings** area, see "ID mapping."
- For details about configuring the settings in the **On-Disk Identity** area, see "On-disk identity selection."

Configure general authentication settings

General settings include the lsass log level, character substitution for spaces encountered in user and group names, and whether to send NTLMv2 responses.

1. On the **File Sharing** menu, point to **Advanced**, and then click **Authentication Settings**.
The Advanced **Authentication Settings** page appears.
2. In the **General Settings** section, configure the following settings as needed.



Note: To configure a setting, you must first click to select its check box. If a setting is already configured, clearing its check box will make that configuration the default setting.

- **Lsass log level:** This setting specifies the default logging level for the authentication daemon (lsassd) across the cluster.

Each successive level from **error** to **trace** provides more detailed information. These logs are primarily designed to be read and interpreted by Isilon Product Support. Acceptable values are: **error** (default), **warning**, **info**, **verbose**, **debug**, **trace**.
- **Send NTLMv2:** When connecting as an SMB client, this setting configures the type of NTLM response that is sent. NTLMv2 provides additional security over NTLM and should be used if all servers support the protocol. Acceptable values are: **yes**, **no** (default).

- **Space replacement:** Some clients have difficulty handling spaces in user and group names. This setting causes lsassd to substitute a character whenever a space is encountered. Care should be taken when choosing a character so that it does not conflict with characters already used by the name.

3. Click **Submit**.

Run the Repair Permissions job

You can use the Repair Permissions job to correct file permissions or identities.

1. On the **File Sharing** menu, point to **Advanced**, and then click **Repair Permissions Job**.
The Advanced **Repair Permissions Job** page appears.
2. For **Repair task**, click to select one of the following settings:
 - **Convert permissions:** For each file and directory within the specified **Path to repair** setting, converts the owner, group and access control list (ACL) to the target on-disk identity. To prevent permissions issues, this task should be run whenever the on-disk identity has been changed.
 - **Clone permissions:** Applies the permissions settings for the specified **Template Directory** as-is to the directory specified in the **Path to repair**.
 - **Inherit permissions:** Recursively applies the ACL that is used by the specified **Template Directory** to each file and subdirectory within the specified **Path to repair** directory, according to normal inheritance rules.
3. To change the priority level of this job compared to other jobs, click a value (**1-10**) in the **Priority** box.
4. To specify a different impact policy for this job to use, click an available option in the **Impact policy** box.
5. For **Path to repair**, type the full path beginning at `/ifs` to the directory whose permissions need repaired, or click **Browse** to navigate to the directory via File System Explorer.
6. For **Template Directory** (available with Clone and Inherit tasks only), type the full path beginning at `/ifs` to the directory whose permissions settings you want to apply, or click **Browse** to navigate to the directory via File System Explorer.
7. For **Target** (available with Convert task only), optionally select the on-disk identity type:
 - **Use default system type:** Uses the system's default identity type. This is the default setting.
 - **Use native type:** If a user or group does not have a real UNIX identifier (UID or GID), uses the Windows identity type (SID).
 - **Use UNIX type:** Uses the UNIX identity type.
 - **Use SID (Windows) type:** Uses the Windows identity type.

Identity management

OneFS supports three primary identity types, each of which can be stored directly on the file system: user identifier (UID) and group identifier (GID) for UNIX, and security identifier (SID) for Windows. These identity types are used when creating files, checking file ownership or group membership, and performing file access checks. In OneFS, names are classified as a secondary identifier and are used for authorization but never for authentication.

UNIX and Windows identifiers are formatted as follows:

- A UID or GID is a 32-bit number (with a maximum of 4,294,967,295) and is the traditional way to refer to users or groups on the system. This is sufficient for UNIX identity sources, but some mapping is required to function correctly with Microsoft Active Directory.
- An SID is a series of authorities and sub-authorities ending with a 32-bit relative identifier (RID). Most SIDs have the form `S-1-5-21-A-B-C-<RID>`, where A, B and C are specific to a domain or computer and `<RID>` denotes the object inside the domain. An SID is the primary identifier for users and groups in Active Directory.

Multiple names can reference the same object within OneFS. There are many variations in the way a name can be entered or displayed:

- UNIX assumes unique case-sensitive namespaces for users and groups (for example, Name and name represent different objects).
- Windows provides a single, case-insensitive name space for all objects and also specifies a prefix to target a specific Active Directory domain (for example, domain\name).
- Kerberos and NFSv4 define principals, which require all names to be formatted similar to email addresses (for example, name@domain).

For example, given the name "support" and the domain EXAMPLE.COM, support, EXAMPLE\support, and support@EXAMPLE.COM are all names for a single object in Active Directory.

Whenever a name is provided as an identifier, it is converted into the corresponding object and the correct identity type requested.

Access tokens

Access tokens form the basis of who you are when performing actions on the cluster, and supply the primary owner and group identities to use during file creation. Access tokens are also compared against the ACL or mode bits during authorization checks.

The access token includes all UIDs, GIDs and SIDs for an identity. OneFS exclusively uses the information in the token when determining if a user has access to a particular resource. It is important that the token contains the correct list of UIDs, GIDs and SIDs at all times.

An access token is created from one of the following sources:

Source	Authorization method
Username	<ul style="list-style-type: none"> • SMB impersonate user • Kerberized NFSv3 • Kerberized NFSv4 • mountd root mapping • HTTP • FTP
Privilege Attribute Certificate (PAC)	<ul style="list-style-type: none"> • SMB NTLM • Active Directory Kerberos
User identifier (UID)	<ul style="list-style-type: none"> • NFS AUTH_SYS mapping

Access token generation

For most protocols, the access token is generated from the provided username or by using the authorization data retrieved during authentication. SSH and console login are exceptions to this rule; for these cases, no mapping lookup is performed and the access token contains only the information found by running the `id` command at the command line.



Note: The resulting access token will not contain any SIDs. If you want to use SSH and console login to manage files and access, you should set the on-disk identity to `unix` mode.

The process of token generation and user mapping is described below:

1. Using the initial identity, the user is looked up in all configured authentication providers in order. The user identity and group list is retrieved from the authenticating provider. Sometimes this comes in the form of a Privilege Attribute Certificate (PAC). Any SIDs, UIDs, or GIDs (if present) are added to the initial token.

2. All identities in the token are queried in the ID mapper. All SIDs are converted to their equivalent UID/GID and vice versa. These ID mappings are also added to the access token.
3. If the username matches any user mapping rules, they are processed in order and the token is updated accordingly. (For details about user mapping rules, see "User mapping.")

The default on-disk identity is calculated using the final token and the global setting. These identities are used for newly created files.

User mapping

User mapping provides a way to control the permissions given to users by specifying which user and group identifiers (SIDs, UIDs, and GIDs) the user has. These identifiers are used when creating files and when checking file ownership or group membership. Mapping rules can be used to rename users, add supplemental user identities, and modify a user's group membership. User mapping is only performed during login.

Create a user mapping rule

User mapping rules map user identifiers between Windows file sharing clients and UNIX file sharing clients during login.



Note: Changes made to the user mapping rules will only take effect on newly created connections.

1. On the **File Sharing** menu, point to **Advanced**, and then click **User Mapping Rules**.
The Advanced **User Mapping Rules** page appears.
2. In the **Mapping Rules** section, click **Modify Mapping Rules**.
The **Modify Mapping Rules** dialog box opens.
3. Under **Mapping Rules**, click **Add**.
The **Operation** field displays.
4. In the **Operation** list, click to select one of the following operations:
 - **Append fields from a user:** Modifies an existing credential by adding specified fields to it. All appended identifiers become members of the additional groups list.
 - **Insert fields from a user:** Modifies an existing credential by adding specified fields from another credential. An inserted primary user/group becomes the new primary user/group in the credential and moves the old primary user/group to the additional identifiers list. Modifying the primary user leaves the credential's username unchanged. When inserting additional groups from a credential, the new groups are added to the existing groups.
 - **Replace a user with a new user:** Replaces the existing credential with the credential identified by another user. If another user is not specified, the credential is removed from the list and no user inserted to replace it. If there are no credentials in the list, access is denied with a No Such User error.
 - **Remove supplemental groups from a user:** Modifies a credential by removing the supplemental groups.
 - **Join together two users:** Inserts the new credential into the list of credentials. If the new credential is the second user, it is inserted after the existing credential; otherwise, it is inserted before the existing credential. This is primarily relevant when the existing credential is already the first in the list because the first credential is used to determine the ownership of new system objects.
5. Fill in the fields as needed.



Note: Available fields differ depending on the selected operation.

6. Click **Save Rule**.
7. If there is more than one rule, optionally change the order in which a rule is processed by clicking and dragging it to another location in the list.



Note: To ensure that each rule gets processed, you should order replacements first and allow/deny rules last. For more information, see "Best practices for multiple mapping rules."

8. Click **OK**.

Modify a user mapping rule

You can modify a mapping rule's configured settings. If more than one rule is configured, you can change the order in which they are processed.

1. On the **File Sharing** menu, point to **Advanced**, and then click **User Mapping Rules**.
The Advanced **User Mapping Rules** page appears.
2. In the **Mapping Rules** section, click **Modify Mapping Rules**.
The **Modify Mapping Rules** dialog box opens.
3. Under **Mapping Rules**, click the rule that you want to modify.
The fields for the current operation are displayed.
4. Modify the settings as needed.
5. Click **Save Rule**.
6. Optionally change the order in which a rule is processed by clicking and dragging it to another location in the list.
7. Click **OK**.

Delete a user mapping rule

Follow this procedure to delete a user mapping rule.

1. On the **File Sharing** menu, point to **Advanced**, and then click **User Mapping Rules**.
The Advanced **User Mapping Rules** page appears.
2. In the **Mapping Rules** section, click **Modify Mapping Rules**.
The **Modify Mapping Rules** dialog box opens.
3. Under **Mapping Rules**, click the rule that you want to delete, and then click **Delete**.
4. Click **OK**.

Test user mapping rules

Before mapping rules are applied, you can test the configured settings to determine the effect that those rules will have. The mapping rule test does not make any changes but allows you to verify that the rules produce the correct identifiers.

1. On the **File Sharing** menu, point to **Advanced**, and then click **User Mapping Rules**.
The Advanced **User Mapping Rules** page appears.
2. In the **Test Mapping** section, in the **User** field, type a username or click **Choose user** to search for a user.
3. Click **Test Mapping**.
The **Test Results** section appears, and displays the user's primary identifiers, on-disk identities, and any additional identities.

Best practices for multiple mapping rules

To minimize complexity when configuring multiple mapping rules, we recommend grouping and organizing rules by type.

By default, every mapping rule is processed so that multiple rules can be applied. This can present problems when applying a “deny all” rule, such as denying any unknown user. Additionally, replacement rules may interact with rules containing wildcards.

We recommend grouping rules by type and ordering the groups as follows:

1. **Replacements.** Any user renaming should be processed first to ensure that all instances of the name are replaced.
2. **Joins.** After the names are set by any replacement operations, use join, add and insert rules to add extra identifiers.
3. **Allow/deny.** All processing must be stopped before a default deny rule can be applied. To do this, create a rule which matches allowed users but does nothing (such as an add operator with no field options specified), and has the break option. After enumerating the allowed users, a catchall deny may be placed at the end to replace anybody unmatched with an empty user.

Within each of these groups, it is best to place explicit rules before rules involving wildcards. Otherwise the wildcard may apply first, causing the explicit rule to be skipped over.

ID mapping

The file access protocols provided by OneFS support a limited number of identity types, usually either UIDs and GIDs or just SIDs. When an identity type is requested that does not match the stored type, a mapping is required.

Mappings are stored in a cluster-distributed database called the ID mapper. When retrieving a mapping from this database, as input the ID mapper takes a source and target identity type (UID, GID, or SID). If a mapping already exists between the specified source and the requested type, that mapping is returned; otherwise, a new mapping is created.

Each mapping is stored as a one-way relationship from source to destination. Two-way mappings are presented as two complementary one-way mappings in the database.

There are four types of identity mappings, two of which are stored persistently in the ID mapper by default.

- *External mappings* are derived from identity sources outside of OneFS. For example, Active Directory (AD) can store a UID or GID along with an SID. When retrieving the SID from AD, the UID/GID is also retrieved and used for mappings on OneFS. By default, mappings derived from AD are not persistently stored in the ID mapper, but mappings from other external identity sources (LDAP, NIS) are persistently stored.
- *Algorithmic mappings* are created by adding a UID or GID to a well-known base SID, resulting in a “UNIX SID.” (For more information, see “UID and GID mappings to SIDs.”) Algorithmic mappings are not persistently stored in the ID mapper database.
- *Manual mappings* are set explicitly by running the `isi auth mapping` command at the command line. For command syntax and examples, see the OneFS Command Reference. Manual mappings are stored persistently in the ID mapper database.
- *Automatic mappings* are generated if no other mapping type can be found. An SID is mapped to a UID or GID out of the default range of 1,000,000-2,000,000. This range is assumed to be otherwise unused and a check is made only to ensure there is no mapping from the given UID before it is used. After creation, these mappings are stored persistently in the ID mapper database.

UID and GID mappings to SIDs

If a UID-to-SID or GID-to-SID mapping is requested, a temporary “UNIX SID” is generated in the format `S-1-22-1-<UID>` or `S-1-22-2-<GID>`.

UIDs and GIDs have a set of pre-defined mappings to and from SIDs.

- UIDs are mapped to an SID with a domain of `S-1-22-1` and a resource ID (RID) matching the UID. For example, the “UNIX SID” for UID 600 is `S-1-22-1-600`.
- GIDs are mapped to an SID with a domain of `S-1-22-2` and a RID matching the GID. For example, the “UNIX SID” for GID 800 is `S-1-22-2-800`.

UID-to-SID and GID-to-SID mappings are only used if the caller requests a mapping to be created and one did not already exist. The resulting “UNIX SIDs” will never be stored on disk.

SID mappings to UIDs and GIDs

If an SID-to-UID or SID-to-GID mapping is requested, the actual object (usually an AD user or group) must first be located.

After the object is located, the following rules are applied to create two mappings, one in each direction:

1. If the object has an associated UID or GID (an external mapping), create a mapping from the SID.
2. If a mapping for the SID already exists in the ID mapper database, use that mapping.
3. Determine whether a lookup of the user or group is necessary in an external source, according to the following conditions:
 - The user/group is in the primary domain or one of the listed lookup domains.
 - Lookup is enabled for users/groups.

4. If a lookup is necessary, follow these steps:
 - a. By default, normalize the user/group name to lowercase.
 - b. Search all sources except Active Directory (that is, local, LDAP, and NIS) for a matching user or group object.
 - c. If an object is found, use the associated UID or GID and create an "external" mapping.
5. Allocate an automatic mapping from the configured range.

Configure settings for ID mapping

Administrators with advanced knowledge of UNIX and Windows identities can modify the default settings that determine how those identities are mapped in the system.

ID mapping parameters include UID and GID allocation, lookup options, and NetBIOS domain mapping for untrusted domains.



Caution: Modifying these settings can cause authentication issues.

1. On the **File Sharing** menu, point to **Advanced**, and then click **Authentication Settings**.
The Advanced **Authentication Settings** page appears.
2. In the **ID Mapping Settings** section, configure the following settings as needed.



Note: To configure a setting, you must first click to select its check box. If a setting is already configured, clearing its check box will make that configuration the default setting.

- **Allocate gids:** Use "automatic" mapping to assign a GID from the default 1,000,000-2,000,000 range if no other SID-to-GID mapping can be found. Acceptable values are: **yes** (default), **no**.
- **Allocate uids:** Use "automatic" mapping to assign a UID from the default 1,000,000-2,000,000 range if no other SID-to-UID mapping can be found. Acceptable values are: **yes** (default), **no**.
- **Lookup domains:** When one or both **Lookup groups** and **Lookup users** settings are configured, this setting specifies the Active Directory (AD) domains for which SIDs will be mapped to UIDs and GIDs via name matching. By default, only accounts in the primary AD domain are looked up.
- **Lookup groups:** Before resorting to allocating an automatic GID for a group, OneFS searches the local, LDAP, and NIS sources for groups that match by name. If such an account is found, its GID is used. Acceptable values are: **yes** (default), **no**.
- **Lookup normalize groups:** This setting affects the behavior of **Lookup groups** by normalizing the group's name to lowercase before searching. This could be undesirable if a source contains case-sensitive names. Acceptable values are: **yes** (default), **no**.
- **Lookup normalize users:** This setting affects the behavior of **Lookup users** by normalizing the user's name to lowercase before searching. This could be undesirable if a source contains case-sensitive names. Acceptable values are: **yes** (default), **no**.
- **Lookup users:** Before resorting to allocating an automatic UID for a user, OneFS searches the local, LDAP, and NIS sources for accounts that match by name. If such an account is found, its UID is used. Acceptable values are: **yes** (default), **no**.
- **Map untrusted:** As part of the NTLMv1 and NTLMv2 protocols, a NetBIOS domain is usually provided by the client computer along with the user name, even if unspecified by the operator. This setting allows lsassd to map unknown domains into the one specified. If no domain is specified, the user is assumed to exist in the local domain and will be authenticated against the local user database. Specifying an AD domain can cause issues with NTLMv2 authentication.

3. Click **Submit**.

On-disk identity selection

Choosing the preferred identity to store on disk is important because nearly all protocols require some level of mapping in order to operate correctly. You can choose to store the UNIX or the Windows identity, or allow the system to determine

the correct identity to store; however, the on-disk selection does not guarantee the preferred identity can always be stored on disk.



Note: When upgrading from a previous version of OneFS, the on-disk identity is set to `unix`. On new installations and re-imaging, the on-disk identity is set to `native`.

If the `unix` on-disk identity type is set, the system authentication daemon (`lsassd`) looks up incoming SIDs in the configured authentication sources. If a match is found, the SID is converted to either a UID or GID. If the identity does not exist on the cluster (for example, it is local to the client or part of an untrusted AD domain), the SID is stored instead.

Similarly, for the `sid` on-disk identity, `lsassd` searches the configured authentication sources for matches to an incoming UID or GID. If no match is found, the UNIX ID is stored on disk.

For the `native` on-disk identity type, `lsassd` attempts to choose the correct identity to store on disk by checking for the following ID mapping types in order:

1. Algorithmic mapping: An SID that matches `S-1-22-1-<UID>` or `S-1-22-2-<GID>` is converted back to the corresponding UNIX ID and set as the on-disk identity.
2. External mapping: An object that has an explicit UNIX ID defined in an external source (AD, LDAP, or NIS) has that ID set as the on-disk identity.
3. Mappings stored persistently in the ID mapper database: An identity with a persistent mapping in the ID mapper database uses the destination of that mapping as the on-disk identity. This pertains primarily to manual mappings. For example, if a mapping of `GID:10000 -> S-1-5-32-545` exists, a request for the on-disk storage of `GID:10000` returns `S-1-5-32-545`.

Configure the on-disk identity

You can configure the OneFS operating system to store the UNIX or Windows identity on-disk, which is especially useful for users of Active Directory. OneFS uses the on-disk identity to transparently map identities at a global level for individual protocols.



Caution: This setting controls how identities are stored in the file system, and should not be changed without fully understanding its potential effect on protocol behavior. Improperly changing this setting can cause compatibility issues for connecting clients, resulting in access-denied or worse errors.

1. On the **File Sharing** menu, point to **Advanced**, and then click **Authentication Settings**.
The Advanced **Authentication Settings** page appears, and displays the **On-Disk Identity** section.
2. Click to select the **On disk identity** check box, and then select one of the following options.
For guidance on selecting the on-disk identity, see "On-disk identity selection."
 - **native:** If a user or group does not have a real UNIX identifier (UID or GID), store its Windows identifier (SID).
 - **unix:** Always store the UNIX identifier, if available. This setting is recommended for NFSv2 and NFSv3, which use UIDs and GIDs exclusively.
 - **sid:** Always store the Windows identifier, if available.
3. Click **Submit**.
The confirmation message, "Successfully saved settings" displays at the top of the page.
4. To prevent permission errors, run the Repair Permissions task 'Convert Permissions.' This task should be run whenever you changes the on-disk identity.
The Convert Permissions task visits all files and directories starting in the specified path and converts the owner, group and ACLs to the target on-disk identity.
 - a. In the **On-Disk Identity** section, click the **Repair Permissions** link.
The Advanced **Repair Permissions Job** page appears.
 - b. Optionally modify the **Priority** and **Impact policy** settings.
 - c. For the **Repair task** setting, click to select **Convert permissions**.
 - d. For the **Path to repair** setting, type or click **Browse** to select the path to the directory whose permissions you want to repair.

- e. For the **Target** setting, ensure the **Use default system type** is selected.
- f. Click **Start**.
The "Successfully started job" message displays at the top of the page.

Authorization

During user authorization, OneFS compares the access token generated during the connection with the authorization data found on the file. All user and identity mapping occurs during token generation; no mapping is performed when evaluating permissions.

Authorization follows one of two different paths depending on whether the file has POSIX mode bits or an ACL.

File authorization data

OneFS supports two types of authorization data on a file: access control lists (ACLs) and UNIX permissions. The type used is based on the ACL policies that are set on the file creation method.

Generally, files that are created over SMB or within a directory that has an ACL will receive an ACL; otherwise, OneFS relies on the POSIX mode bits that define UNIX permissions. In either case, the owner can be represented by a UNIX identifier (UID or GID), or by its Windows identifier (SID). The group can be represented by a GID or SID. Although mode bits are present when a file has an ACL, those bits are provided only for protocol compatibility and are not used for access checks.

When performing an authorization check, OneFS compares the access token generated during the connection with the authorization data found on the file. OneFS does not map identities when evaluating permissions; all mapping takes place during access token generation. If required to evaluate a UNIX permission against a file with an ACL, OneFS converts the permissions into the corresponding rights that the caller must possess.

ACLs and UNIX-style permissions

OneFS is optimized for a mixed environment in which any file or directory can be governed by either a Windows access control list (ACL) or UNIX permissions.

Regardless of the security model, access rights are enforced consistently across access protocols; that is, a user is granted (or denied) the same rights to a file when using SMB (Windows file sharing) as they would when using NFS (UNIX file sharing). Clusters running OneFS support a set of global policy settings that enable you to customize the default ACL and UNIX permissions settings to best support your environment.

By default, OneFS ships with traditional UNIX permissions on the file tree. By using Windows Explorer or OneFS administrative tools, any file or directory can be given an ACL. In addition to Windows domain users and groups, OneFS ACLs can also include local, NIS, and LDAP users and groups. After a file has been given an ACL, its previous mode bits are no longer enforced, and exist only as an estimation of the effective permissions.

By default, OneFS is configured with the optimal settings for a mixed UNIX and Windows environment; however, you can configure ACL policies if necessary to optimize for UNIX or Windows. You should only configure ACL and UNIX permissions, however, if you fully understand how they interact with and affect one another.

POSIX Mode Bits (UNIX permissions)

In a UNIX environment, file and directory access is controlled by POSIX mode bits, which allow read, write, or execute permissions to the owning user, the owning group, and "everyone else."



Note: OneFS supports the standard UNIX tools for changing permissions, `chmod` and `chown`. For more information about these tools, refer to the command-line man pages for the `chmod`, `chown`, and `ls` commands.

All files contain 16 bits, which provide information about the file or directory type and the permissions. The lower 9 bits are grouped as three 3-bit sets, called triples, which contain the read (r), write (w), and execute (x) permissions for each class of users (owner, group, other). You can set permissions flags to grant permissions to each of these classes.

Assuming the user is not root, the class is used to determine if the requested access to the file should be granted or denied. The classes are not cumulative; the first class matched is used. It is therefore common practice to grant permissions in decreasing order.

Windows-style ACLs

In Windows environments, file and directory permissions, referred to as access rights, are defined in access control lists (ACLs). ACLs are more complex than mode bits, but are also capable of expressing much richer sets of access rules. OneFS uses the ACL processing rules commonly associated with Windows ACLs.

A Windows ACL is composed of one or more access control entries (ACEs), each representing the security identifier (SID) of a user or a group as a trustee. In OneFS, an ACL can contain ACEs with a UID, GID, or SID as the trustee. Each ACE in the ACL contains its own set of rights that allow or deny access to a file or folder, and can optionally contain inheritance flags to specify that the ACE should be inherited by any child folders and files.

Instead of the standard 3 permissions available for mode bits, ACLs have 32 bits of fine grained access rights. Of these, the upper 16 bits are general and apply to all object types. The lower 16 bits vary between files and directories but are defined in a compatible way that allows most applications to use the same bits for files and directories.


Rights can be used for granting or denying access for a given trustee. Access can be blocked to a user explicitly through the use of a deny ACE, or implicitly by ensuring that the user does not directly (or indirectly through a group) appear in an ACE that grants the right in question.

Configure basic permissions settings

You can change the Isilon cluster's default access control list (ACL) settings globally or individually, to best support your environment. These global permissions policies change the behavior of permissions on the system.



Note: To change policy settings, you must have advanced knowledge of Windows ACLs. The default ACL settings are sufficient for most cluster deployments and should be modified only as necessary by experienced administrators.

1. On the **File Sharing** menu, point to **Advanced**, then click **ACL Policies**.
The **Advanced > ACL Policies** page appears.
2. In the **Standard Settings** section, under **Environment**, click to enable one of the following options:
 - To cause cluster permissions to operate with UNIX semantics, as opposed to Windows semantics, click **UNIX only**. By enabling this option, you prevent ACL creation on the system.
 - To cause cluster permissions to operate in a mixed UNIX and Windows environment, click **Balanced**. This setting is recommended for most cluster deployments.
 - To cause cluster permissions to operate with Windows semantics, as opposed to UNIX semantics, click **Windows only**.
-  **Note:** Enabling this option causes the system to return an error on UNIX `chmod` requests.
- To configure individual permissions policy settings, click **Configure permission policies manually**.
3. If you enabled **UNIX only**, **Balanced**, or **Windows only**, the corresponding options in the **Permission Policies** section are automatically enabled or disabled as appropriate. Click **Submit** to set your permissions on the cluster.
4. If you enabled **Configure permission policies manually**, in the **Permission Policies** section, configure the individual permission policy settings as described below, and then click **Submit**.
 - To disable ACL creation on the cluster, next to the **ACL creation over SMB** option, click **Do not allow the creation of ACLs over Windows File Sharing (SMB)**. This setting has the greatest impact on Windows file sharing (SMB) clients because the cluster silently fails any operation that sets ACLs. However, it also affects cluster command-line interface commands, such as `chmod`, which you can use to set ACLs.



Note: Inheritable ACLs on the system take precedence over this setting; if inheritable ACLs are set on a folder, any new files and folders created in that folder will inherit the folder's ACL. Disabling this setting

does not remove ACLs currently set on files. To clear an existing ACL, you must run the `chmod` command to set the correct permissions.

- To control what happens when a `chmod` operation is initiated on a file with an ACL, either over UNIX file sharing (NFS) or locally, click **chmod on files with existing ACLs**. Enabling this setting controls any elements that set UNIX permissions, including the cluster's File System Explorer. Enabling this policy setting does not change how `chmod` operations affect files that do not have ACLs.

To cause Windows and UNIX permissions to operate smoothly in a balanced environment, click **Merge the new permissions with the existing ACL**. Enabling this setting causes the `chmod` permissions to merge with already existing ACLs.

An ACE for each identity (owner, group, and everyone) is either modified or created, but all other ACEs are unmodified. Inheritable ACEs are also left unmodified to enable Windows users to continue to inherit appropriate permissions. However, UNIX users can set specific permissions for each of those three standard identities.

If you do not need permissions to be set from Windows, click **Remove the existing ACL and set UNIX permissions instead**. For `chmod` operations, enabling the setting removes any existing ACL and instead sets the `chmod` permissions.

To store the UNIX permissions in a Windows ACL, click **Remove the existing ACL and create an ACL equivalent to the UNIX permissions**. Enable this setting only if you want to remove Windows permissions but do not want files to have synthetic ACLs (for a definition of this type of ACL, see the *Treatment of "rwx" permissions* section below).

To store the UNIX permissions in a Windows ACL, and assign any group rights specified in the new `chmod` operation to users/groups in the old ACL, click **Remove the existing ACL and create an ACL equivalent to the UNIX permissions, for all users/groups referenced in old ACL**.

To prevent users from making NFS and local `chmod` operations, click **Deny permission to modify the ACL**. Enable this setting if you do not want to allow permission sets over NFS.



Note: If you try to run the `chmod` command on the same permissions that are currently set on a file with an ACL, you may cause the operation to silently fail. (The operation appears to be successful, but if you were to examine the permissions on the cluster, you would notice that the `chmod` command had no effect.)

As a workaround, you can run the `chmod` command away from the current permissions and then perform a second `chmod` command to revert to the original permissions. For example, if your file shows 755 UNIX permissions and you want to confirm this number, you could run `chmod 700 file; chmod 755 file`.

- To change a file or folder's owning user or group, click **chown on files with existing ACLs**. Over NFS, the `chown` operation changes the permissions and the owner or owning group. For example, consider a file owned by user Joe with "rwx-----" (700) permissions, signifying "rwx" permissions for the owner, but no permissions for anyone else. If you run the `chown` command to change ownership of the file to user Bob, the owner permissions are still "rwx" but they now represent the permissions for Bob, rather than for Joe. In fact, Joe will have lost all of his permissions.

To cause the `chown` operation to perform as it does in UNIX, click **Modify the owner and/or group permissions**. Enabling this setting modifies any ACEs in the ACL associated with the old and new owner or group.

To cause the NFS `chown` operation to function as it does in Windows, click **Do not modify the ACL**. When a file owner is changed over Windows, no permissions in the ACL are changed.



Note: This setting does not affect UNIX `chown` operations performed on files with UNIX permissions, and it does not affect Windows `chown` operations, which do not change any permissions.

- **Access checks (chmod, chown):** In UNIX environments, only the file owner or superuser has the right to run a `chmod` or `chown` operation on a file. In Windows environments, you can implement this policy setting to give users the right to perform `chmod` operations, called the "change permissions" right, or the right to perform `chown` operations, called the "take ownership" right.



Note: The "take ownership" right only gives users the ability to take file ownership, not to give ownership away.

To cause the `chmod` and `chown` access checks to operate with UNIX-like behavior, click **Allow only owners to chmod or chown**.

To cause the `chmod` and `chown` access checks to operate with Windows-like behavior, click **Allow owner and users with 'take ownership' right to chown, and owner and users with 'change permissions' right to chmod**.

- **Treatment of "rwx" permissions:** In UNIX environments, "rwx" permissions signify two things: A user or group has read, write, and execute permissions; and a user or group has the maximum possible level of permissions.

When you assign UNIX permissions to a file, no ACLs are stored for that file. However, a Windows system processes only ACLs; Windows does not process UNIX permissions. Therefore, when you view a file's permissions on a Windows system, the Isilon cluster must translate the UNIX permissions into an ACL. In the Isilon cluster, this type of ACL is called a *synthetic ACL*. Synthetic ACLs are not stored anywhere; instead, they are dynamically generated as needed and then they are discarded. If a file has UNIX permissions, you may notice synthetic ACLs when you run the `ls -le` command on the cluster in order to view a file's ACLs.

When you generate a synthetic ACL, the cluster maps UNIX permissions to Windows rights. Windows supports a more granular permissions model than UNIX does, and it specifies rights that cannot easily be mapped from UNIX permissions. If the cluster maps "rwx" permissions to Windows rights, you must enable one of the following options: The **Retain 'rwx' permissions** setting generates an ACE that gives read, write, and execute permissions, and the setting **Treat 'rwx' permissions as Full Control** generates an ACE that provides the maximum Windows permissions for a user or a group.



Note: The main difference between "rwx" and "Full Control" is the broader set of permissions. "Full control" adds the "change permissions" right, the "take ownership" right, and the "delete" right.

To enable the synthetic ACL to provide only read, write, and execute permissions, click **Retain 'rwx' permissions**.

To give the synthetic ACL *Full Control*, click **Treat 'rwx' permissions as Full Control**.

Configure advanced permissions settings

Advanced access control list (ACL) and permissions settings, which you configure separately from the basic settings, are not limited by the environment. Instead, you can apply these settings to all environment types supported by the Isilon cluster including UNIX, Windows, and balanced.



Note: To change policy settings, you must have advanced knowledge of Windows ACLs.

1. On the **File Sharing** menu, point to **Advanced**, then click **ACL Policies**.
The **Advanced > ACL Policies** page appears.
2. In the **Advanced Settings** section, enable the **Group owner inheritance** policy if you want to assign a group owner of new files and folders.



Note: Operating systems tend to work with group ownership and permissions in two different ways: BSD inherits the group owner from the file's parent folder. Windows and Linux inherit the group owner from the file creator's primary group.

- To control file behavior based on whether the new file inherits ACLs from its parent folder, click **When an ACL exists, use Linux and Windows semantics, otherwise use BSD semantics** (the default setting). If it does, the file uses the creator's primary group. If it does not, the file inherits from its parent folder.
- To cause the group owner to be inherited from the file's parent folder, click **BSD semantics - Inherit group owner from the parent folder**.
- To cause the group owner to be inherited from the file creator's primary group, click **Linux and Windows semantics - Inherit group owner from the creator's primary group**.



Note: If you enable a setting that causes the group owner to be inherited from the creator's primary group, it can be overridden on a per-folder basis by running the `chmod` command to set the **set-gid** bit. For more information, see the manual page for the `chmod` command.

3. Configure the **chmod (007) on files with existing ACLs** policy setting.

- To set 007 UNIX permissions without removing an existing ACL, click **chmod(007) does not remove existing ACL**.
- To remove ACLs from over UNIX file sharing (NFS) and locally on the cluster through the `chmod (007)` command, click **chmod(007) removes existing ACL and sets 007 UNIX permissions**.



Note: If you enable the **chmod (007) on files with existing ACLs** policy setting, be sure to run the `chmod` command on the file immediately after using `chmod (007)` to clear an ACL. In most cases, you do not want to leave 007 permissions on the file.

4. Set the **Owner permissions and group permissions** policies: It is impossible to represent the breadth of a Windows ACL's access rules using a set of UNIX permissions. Therefore, when a UNIX client requests UNIX permissions for a file with an ACL over NFS (an action known as a "stat"), it receives an imperfect approximation of the file's true permissions. By default, executing an `ls -l` command from a UNIX client returns a more open set of permissions than the user expects. This permissiveness compensates for applications that incorrectly inspect the UNIX permissions themselves when determining whether to attempt a file-system operation. The purpose of this policy setting is to ensure that these applications proceed with the operation to allow the file system to properly determine user access through the ACL.

- To make the owner or group permissions appear more permissive than the actual permissions on the file, click **Approximate owner [or group] mode bits using all possible group ACEs with the owner ID**.
- To make the owner or group permissions appear more accurate, in that you see only the permissions for a particular owner or group and not the more permissive set, click **Approximate owner [or group] mode bits using only the ACE with the owner ID**. However, this may cause access-denied problems for UNIX clients.

5. Set the **No "deny" ACEs** policies: The Windows ACL user interface cannot display an ACL if any "deny" ACEs are out of canonical ACL order. However, in order to correctly represent UNIX permissions, deny ACEs may be required to be out of canonical ACL order.

- Enabling **Do not modify synthetic ACLs** does not modify synthetic ACL generation; "deny" ACEs will be generated when necessary and may be out of order.
- To initiate the removal of any "deny" ACEs when generating synthetic ACLs, click **Remove "deny" ACEs from synthetic ACLs**.



Note: The **Remove "deny" ACEs from synthetic ACLs** option causes permissions to be incorrectly approximated. This incorrect approximation is permanently set if a Windows user or an application performs an ACL get, an ACL modification, and an ACL set (known as a "roundtrip") to and from Windows.

6. Click **Submit**.

Authorization data mapping in mixed environments

Whenever a file operation requests an object's authorization data—for example, via the `ls -l` command over NFS or the **Security** tab over SMB—OneFS attempts to provide that data in the requested format. In a mixed UNIX/Windows environment, some mapping may be required when performing create file, set security, get security, or access operations.

Windows ACL mapping to UNIX permissions

If the file contains an owner or group that is an SID, the system attempts to map it into a corresponding UID or GID before returning it.

In UNIX, authorization data is retrieved by calling `stat(2)` on a file and looking at the owner, group, and mode bits. Over NFSv3, the `GETATTR` command functions similarly. The mode bit approximation is calculated and set on the file whenever its ACL changes, and only needs to be retrieved to service these calls.



Note: Because the resulting SID-to-UID or -GID mappings are cached in both the ID mapper and the `stat` cache for easy retrieval, if the mapping has recently changed the file could report inaccurate information until the file is updated or the cache is flushed via the `isi_flush` command.

UNIX permissions mapping to Windows ACL

No UID/GID-to-SID mappings are performed when creating an ACL for a file; all UIDs and GIDs are converted to SIDs (or principals, for NFSv4) when the ACL is returned.

OneFS uses a two-step process for returning a security descriptor, which contains SIDs for the owner and primary group of an object:

1. The existing security descriptor is retrieved from the file. If the file does not have a DACL (discretionary access control list), a synthetic ACL is constructed from the file's lower 9 mode bits, which are separated into three sets of permission triples (one each for owner, group, and everyone). For more details about mode bits, see "POSIX Mode Bits (UNIX permissions)."
2. Two access control entries (ACEs) are created for each triple: the Allow ACE contains the corresponding rights that are granted according to the permissions; the Deny ACE contains the correspond rights that are denied. In both cases the trustee of the ACE corresponds to the file owner, group, or everyone. After all of the ACEs are generated, any that are not needed are stripped out before the synthetic ACL is returned.

Job management

The OneFS operating system includes a number of jobs that you can control and modify. You cannot create or delete jobs.

You can set or change the priority level of a job in relation to other jobs, as well as the impact policy. For jobs that are not automatically scheduled, you can modify the interval and frequency for running the job.

Monitor jobs

You can view information about running, paused, waiting, and failed jobs.

1. On the **Cluster** menu, point to **Operations**, and then click **Summary**.

The Operations **Summary** page appears, and displays the following sections:

- **Running Jobs:** Displays any jobs that are currently running on the cluster. You can update, pause, or cancel running jobs.
- **Paused and Waiting Jobs:** Displays any jobs that are queued to start, or that have been paused. You can update, resume, or cancel paused jobs, and you can update or cancel waiting jobs.
- **Jobs to Retry:** Displays any failed jobs. You can cancel or retry failed jobs.

Depending on the state of the job, the following information is displayed:

- **Status:** Displays the current status of the job. Running jobs are indicated by a rotating Busy symbol. Other possible job states are indicated by color-coded icons and include User Paused (yellow), Waiting (blue), or Failed (red).
- **Job ID:** Displays the job's ID number.
- **Job:** Displays the name of the job. You can view a job's details by clicking its name.
- **Priority:** Displays the job's assigned priority level.
- **Impact Policy:** Displays the job's impact policy.
- **Elapsed:** Displays the job's elapsed duration.
- **Phase:** Displays the current phase of the job.
- **Progress:** Displays the job's progress.
- **Ended (Failed jobs only):** Displays the timestamp for when the job ended.
- **Retries (Failed jobs only):** Displays the number of times the job will be retried.
- **Errors (Failed jobs only):** Displays the number of errors that the job encountered.

2. Optionally, click a job name to view a comprehensive summary of the job's configuration settings, current state, and recent job events.

Modify a job

You can modify a job's priority level, impact policy, and schedule settings. You cannot modify a job's name or description.

1. On the **Cluster** menu, point to **Operations**, and then click **Jobs and Impact Policies**.
The Operations **Jobs and Impact Policies** page appears.
2. In the **Available Jobs** section, under **Actions**, click **Edit** for the job that you want to modify.

The Operations **Modify Job** page appears.

3. To change the priority level of this job compared to other jobs, click a value (**1-10**) in the **Priority** box.
4. To specify a different impact policy for this job to use, click an available option in the **Impact policy** box.
5. To add or modify the schedule for this job, click **Edit schedule** and then, in the **Edit Schedule** dialog box, modify the interval and frequency settings as needed.



Note:

- This option is not available for automatically scheduled jobs.
 - To modify the AVScan job schedule, you must use the Anti-Virus **Policies** page.
- a. To change the schedule's interval, select a new value (**Daily, Weekly, Monthly, or Yearly**). You can optionally change the default settings.
 - b. To change the schedule's frequency, select a new value (**Once or Multiple times**). You can optionally change the default settings.
 - c. Click **Done**.
6. Click **Submit**.

Start a job

In addition to scheduling jobs, you can start a job at any time.

This procedure describes how to start a job from the Operations **Jobs and Impact Policies** page. You can also start a job from the **Summary** page or the **View Job** page.

1. On the **Cluster** menu, point to **Operations**, and then click **Jobs and Impact Policies**.
The Operations **Jobs and Impact Policies** page appears. The **Available Jobs** section displays the following jobs:
 - **AutoBalance**: Balances free space in the cluster.
 - **AVScan**: Scans files for viruses.
 - **Collect**: Reclaims disk space that could not be freed due to node or drive unavailability.
 - **FlexProtect**: Reprotects the file system.
 - **FSAnalyze**: Gathers file system analytics data.
 - **IntegrityScan**: Verifies file system integrity.
 - **MediaScan**: Scrubs disks for media-level errors.
 - **MultiScan**: Runs Collect and AutoBalance jobs concurrently.
 - **QuotaScan**: Updates quota accounting for domains created on an existing file tree.
 - **SetProtectPlus**: Applies the default file policy. This job is disabled if SmartPools is activated on the cluster.
 - **SmartPools**: Enforces SmartPools file policies.
 - **SnapshotDelete**: Frees disk space that is associated with deleted snapshots.
 - **TreeDelete**: Deletes a path in the `/ifs` directory.
2. Optionally, in the **Impact Policies** section, add, copy, or modify an impact policy for the job. For more information, see "Impact policy management."
3. In the **Available Jobs** section, under **Actions**, click **Start** for the job that you want to start.
The Operations **Start Job** page appears.
4. Optionally update the job's priority level and impact policy.
5. Click **Start**.

Control a job

Depending on the current state of a job, you can pause, resume, cancel, or retry the job. Pausing, resuming, or canceling a job does not affect the state of its associated policy.

1. On the **Cluster** menu, point to **Operations**, and then click **Summary**.
The Operations **Summary** page appears.
2. Perform one of the following actions to change the state of the job:
 - To cancel a running, paused, waiting, or failed job, under **Actions**, click **Cancel**. After performing this action, the job is removed from the Operations **Summary** page.
 - To pause a running job, under **Actions**, click **Pause**. After performing this action, the job is moved to the **Paused and Waiting Jobs** section.
 - To resume a paused job, under **Actions**, click **Resume**. After performing this action, the job is moved to the **Paused and Waiting Jobs** section.
 - To retry a failed job, under **Actions**, click **Retry**.

Update a job

You can modify the priority level and impact policy of a running, paused, or waiting job.

1. On the **Cluster** menu, point to **Operations**, and then click **Summary**.
The Operations **Summary** page appears.
2. In the **Running Jobs** or **Paused and Waiting Jobs** section, under **Actions**, click **Update** for the job whose settings you want to modify.
The Operations **Update Job** page appears.
3. Modify the settings as needed.
4. Click **Update**.

View job history

You can view the history of jobs that have recently been run on the cluster.

1. On the **Cluster** menu, point to **Operations**, and then click **Job History**.
The Operations **Job History** page appears, and the following event information is displayed for each recent job that has been run:
 - **Timestamp**: Displays the time and date of the event.
 - **ID**: Displays the unique identification number of the job instance associated with the event.
 - **Job**: Displays the name of the job associated with the event.
 - **State**: Displays the job state for the event.
 - **Message**: Indicates the beginning or end of a phase, or displays the job state for the event.
 - **Results**: At the end of a phase, displays a summary of that phase.
2. Optionally, click a job name to view a comprehensive summary of the job's configuration settings, current state, and recent job events.

Impact policy management

Impact policies allow you to allocate cluster resources for jobs.

An impact policy includes one or more *impact intervals*, which define a block of time within a given week. Each impact interval is configured to use a single pre-defined *impact level*, which specifies the amount of cluster resources to use.

The following impact levels are available:

- **Paused:** Do not use cluster resources.
- **Low:** Use 10% or less of cluster resources.
- **Medium:** Use 30% or less of cluster resources.
- **High:** Use unlimited cluster resources.

View impact policy settings

You can view information about the impact policies that are configured for use by available jobs.

1. On the **Cluster** menu, point to **Operations**, and then click **Jobs and Impact Policies**.
The Operations **Jobs and Impact Policies** page appears.
2. In the **Impact Policies** section, review the configured impact policies.

Each policy includes the following settings:

- **Name:** Displays the name of the impact policy.
- **Description:** Displays a description of the impact policy, if one exists.
- **Impact Schedule:** Displays the impact level for each configured interval of time when a job that uses this policy is scheduled to run.

By default, OneFS includes the following impact policies. These impact policies cannot be modified or deleted:

- **LOW:** Always use the Low impact level (10% or less of cluster resources).
- **MEDIUM:** Always use the Medium impact level (30% or less of cluster resources).
- **HIGH:** Always use the High impact level (unlimited cluster resources).
- **OFF_HOURS:** Pause the job during business hours, and use the Low impact level during non-business hours.

Create an impact policy

You can create and configure new impact policies for use by scheduled jobs.

This procedure describes how to create an impact policy from scratch. Alternately, you can copy an existing impact policy and modify its settings. For more information, see "Copy an impact policy."

1. On the **Cluster** menu, point to **Operations**, and then click **Jobs and Impact Policies**.
The Operations **Jobs and Impact Policies** page appears.
2. In the **Impact Policies** section, click **Add impact policy**.
The Operations **Add Impact Policy** page appears.
3. In the **Name** box, type a name for the impact policy. This field is required.
4. In the **Description** box, type an optional description of the impact policy.
5. Click **Submit**.
The Operations **Edit Impact Policy** page appears, and displays the **Impact Schedule** section.
6. To modify the default interval settings for this impact policy, click **Edit** and then, on the Operations **Edit Impact Interval** page, configure the following settings as needed:
 - **Start:** Optionally specify the day of the week, the hour, and the minute for the impact interval to start.

- **End:** Optionally specify the day of the week, the hour, and the minute for the impact interval to end.
 - **Impact:** Optionally specify the impact level (**Paused**, **Low**, **Medium**, or **High**) to use for the duration of the impact interval.
7. To add an impact interval for this policy, click **Add impact interval** and then, on the **Add** page, configure the settings as needed.



Note: Any existing impact intervals that are overlapped by the new interval will be overwritten.

8. Click **Submit**.

Copy an impact policy

You can create an exact, configurable copy an existing impact policy.

1. On the **Cluster** menu, point to **Operations**, and then click **Jobs and Impact Policies**.
The Operations **Jobs and Impact Policies** page appears.
2. In the **Impact Policies** section, under **Actions**, click **Copy** for the impact policy that you want to copy.
A new impact policy is created. The new policy name includes the name of the originating policy, appended with `_copy`. For example, if you copy the LOW impact policy, the new policy is named LOW_copy by default.
3. Optionally modify the policy name, description, and impact schedule as needed.

Modify an impact policy

You can modify an impact policy's name, description, and impact schedule.

1. On the **Cluster** menu, point to **Operations**, and then click **Jobs and Impact Policies**.
The Operations **Jobs and Impact Policies** page appears.
2. In the **Impact Policies** section, under **Actions**, click **Edit** for the impact policy that you want to modify.
The Operations **Edit Impact Policy** page appears.
3. Modify the settings as needed.
4. Click **Submit**.

Delete an impact policy

You can remove an impact policy if it is no longer needed.

1. On the **Cluster** menu, point to **Operations**, and then click **Jobs and Impact Policies**.
The Operations **Jobs and Impact Policies** page appears.
2. In the **Impact Policies** section, under **Actions**, click **Delete** for the impact policy that you want to remove.
A confirmation dialog box appears.
3. Click **Yes**.

SmartPools

SmartPools enables you to define subgroups of nodes, called disk pools, within a single OneFS cluster. SmartPools includes an optional license that unlocks additional features including the ability to associate individual files and directories to specific disk pools.

For additional information about SmartPools, or to activate SmartPools for your Isilon IQ clustered storage system, contact your Isilon Systems sales representative.

SmartPools overview

SmartPools applies the information lifecycle management (ILM) concept to disk pools, allowing you to store and move data according to file attributes.

SmartPools includes the following basic features, available with or without a SmartPools license:

- **Disk pools:** Disk pools are dynamic groups of disks associated in a single pool of storage. Disk pool membership changes through the addition or removal of nodes and drives.
- **Virtual hot spare:** You can reserve space in a disk pool, equivalent to up to four full drives, that can be used for data reProtection in the event of a drive failure.



Note: If SmartPools is unlicensed, files are allocated among all available disk pools.

The licensed version of SmartPools includes the following additional features:

- **Custom file pools:** If SmartPools is licensed, you can create custom file pool policies that you can use to filter files and directories into specific disk pools according to attributes such as file size, file type, location, and file creation, change, modification, and access times. The licensed version of SmartPools also includes customizable template policies that are optimized for archiving, extra protection, performance, and VMware files. File pool membership is dynamic, and can change through time or through standard NFS and SMB activity.
- **Disk pool spillover management:** If SmartPools is licensed, you can enable or disable disk pool spillover to define how to handle write operations to a full disk pool. If spillover is enabled, data is redirected to another disk pool; if disabled, new data writes fail. If SmartPools is unlicensed, spillover is forcibly enabled.

SmartPools monitoring

The SmartPools **Summary** page allows you to monitor disk pool usage and view or modify disk pool and file pool policy settings.

1. On the **File System** menu, point to **SmartPools**, and then click **Summary**.
The SmartPools **Summary** page appears.
2. Review the following sections:
 - **Disk Pool Usage:** Displays a chart showing the percentage of available hard disk drive (HDD) and solid-state drive (SSD) space that is currently being used per disk pool.
 - **Disk Pools:** Displays a summary of each disk pool on the cluster, including its provisioning status, resources in the pool, HDD and SSD usage, and default protection level. You can view or modify individual disk pools by clicking **Modify disk pool settings**.

- **File Pool Policies:** Displays a summary of each file pool policy that is configured on the cluster, including an optional description, configured filter and operation settings, and whether any files match the filter criteria. You can view or modify individual file pool policies by clicking **Modify file pool policy settings and rules**.



Note: Disk pool information can also be viewed at the command line by running the `isi status -d -q` command.

SmartPools configuration

You can configure a cluster's directory-level protection policy, global namespace acceleration, virtual hot spare, disk pool spillover, and default protection and I/O optimization settings.



Note: The default protection and default I/O optimization settings comprise the default file pool policy.

Configure basic SmartPools settings

Basic SmartPools settings include directory protection, global namespace acceleration, virtual hot spare, disk pool spillover, protection management, and I/O optimization management.

1. On the **File System** menu, point to **SmartPools**, and then click **Settings**.
The SmartPools **Settings** page appears.
2. In the **SmartPools Settings** section, configure the following settings as needed.
 - **Directory protection:** Increases the amount of protection for directories. To protect directories at a higher level than the directories and files they contain, select the **Protect directories at one level higher** option.
 - **Global namespace acceleration:** Specifies whether to allow per-file metadata to use SSDs in any disk pool.



Note: This setting is available only if more than 20% of the nodes in the cluster contain SSDs and at least 1% of the total cluster storage is SSD-based.

- To restrict per-file metadata to the target pool of the file (except in the case of spillover), click **Disabled**. This is the default setting.
- To allow per-file metadata to use the SSDs in any disk pool, click **Enabled**.
- **Virtual hot spare:** Reserves a minimum amount of space in the disk pool that can be used for data migration in the event of a drive failure.

To reserve disk space for use as a virtual hot spare, select one or both of the following options:

- **Reduce amount of available space:** Subtracts the space reserved for virtual hot spare when calculating available free space.



Note: If this setting is enabled and **Deny new data writes** is disabled, it is possible for the file system utilization to be reported at over 100%.

- **Deny new data writes:** Prevents write operations from using reserved disk space.

Next, configure the **VHS space to reserve** setting. You can reserve a minimum number of virtual drives (1-4), as well as a minimum percentage of total disk space (0-20%). If you configure both settings, the enforced minimum value will satisfy both requirements.

- **Disk pool spillover:** Specifies how to handle write operations to a disk pool that is full.
 - To redirect write operations from a full disk pool to another disk pool, click **Enabled**.
 - To return a disk space error for write operations to a full disk pool, click **Disabled**.



Note: The following options specify whether protection and I/O optimization management through SmartPools is enabled or disabled. Disabling both options effectively disables SmartPools.

- **Protection management:** To use SmartPools to manage disk pool protection, ensure that the **SmartPools manages protection settings** check box is selected. You can optionally modify the default settings under **Default Protection Settings**.



Note: You can overwrite any protection settings that were configured via File System Explorer or the command-line interface by checking the **Including files with manually-managed protection settings** box.

- **I/O optimization management:** To use SmartPools to manage disk pool I/O optimization, ensure that the **SmartPools manages I/O optimization settings** check box is selected. You can optionally modify the default settings under **Default I/O Optimization Settings**.



Note: You can overwrite any I/O optimization settings that were configured via File System Explorer or the command-line interface by checking the **Including files with manually-managed I/O optimization settings** box.

3. Click **Submit**.

Configure default protection settings

You can specify the data pool, snapshot pool, protection level, and SSD strategy for files that are filtered by the default file pool policy.

1. On the **File System** menu, point to **SmartPools**, and then click **Settings**.
The SmartPools **Settings** page appears.
2. In the **SmartPools Settings** section, for the **Protection management** setting, ensure that the **SmartPools manages protection settings** option is selected.



Note: To allow SmartPools to overwrite protection settings that were configured using File System Explorer or the `isi set` command, select the **Including files with manually-managed protection settings** option.

3. In the **Default File Pool Policy Settings** section, under **Default Protection Settings**, modify the settings as needed.
 - To specify the data and snapshot pools, click the **Set data pool to** check box, and then configure the following settings for the data and snapshot pools.
 1. Select **any pool** to assign filtered files to disk pools without restriction, or manually select a pool from the list. For the snapshot pool, you can select **same as data pool**.
 2. If the pool you selected contains SSDs, select one of the following SSD strategies:
 - **Metadata acceleration:** Creates a mirror backup of file metadata on an SSD and writes the rest of the metadata plus all user data on HDDs. Depending on the global namespace acceleration setting, the SSD mirror may be an extra mirror in addition to the number required to satisfy the protection level.
 - **Avoid SSDs:** Never uses SSDs; writes all associated file data and metadata to HDDs only.
 - **Data on SSDs:** Similar to metadata acceleration, but also writes one copy of the file's user data (if mirrored) or all of the data (if not mirrored) on SSDs. Regardless of whether global namespace acceleration is enabled, any SSD blocks reside on the file's target pool if there is room. This SSD strategy does not result in the creation of additional mirrors beyond the normal protection level.
 - To specify the protection level, click the **Set protection level to** check box, and then select one of the following options:
 - **Default protection level of disk pool:** Assigns the default protection policy of the disk pool to the filtered files.

- **Specific level:** Assigns a specified protection policy to the filtered files. To change the protection policy, select a new value from the list.

4. Click **Submit**.

Configure default I/O optimization settings

You can manage the I/O optimization settings that are used in the default file pool policy, including files with manually-managed attributes.

1. On the **File System** menu, point to **SmartPools**, and then click **Settings**.
The SmartPools **Settings** page appears.
2. In the **SmartPools Settings** section, for the **I/O optimization management** setting, ensure that the **SmartPools manages I/O optimization settings** option is selected.



Note: To allow SmartPools to overwrite optimization settings that were configured using File System Explorer or the `isi set` command, select the **Including files with manually-managed I/O optimization settings** option.

3. In the **Default File Pool Policy Settings** section, under **Default I/O Optimization Settings**, select one of the **SmartCache** options.

- To enable SmartCache, click **Enabled**.



Note:

SmartCache can accelerate the process of writing content to the cluster. However, writes to the cluster using SMB will not be affected by SmartCache when write through is specified by the client. Also, writes to the cluster using NFS will not be affected by SmartCache if they are tagged as stable.

If a node crashes with SmartCache enabled while using NFS, unstable writes that have not been committed will be temporarily lost on the node. However, the uncommitted data will be sent to the cluster again as soon as the client reconnects to the cluster, unless the client crashes. As long as the client does not crash before it reconnects, no data will be lost.

- To disable SmartCache, click **Disabled**.

4. Select one of the **Data access pattern** options.
 - To select a random access pattern, click **Random**. This is the default setting for logical units.
 - To select a concurrent access pattern, click **Concurrency**. This is the default setting, with the exception of logical units.
 - To select a streaming access pattern, click **Streaming**. Streaming access patterns can improve performance in some workflows.
5. Click **Submit**.

Reprovision the cluster

Reprovisioning resets all manually provisioned nodes on the cluster to their automatically-provisioned disk pools.

1. On the **File System** menu, point to **SmartPools**, and then click **Settings**.
The SmartPools **Settings** page appears.
2. In the **Default File Pool Policy Settings** section, under **Reprovision Cluster**, click the **Reprovision Cluster** button.
The **Reprovision Cluster** dialog box appears, displaying the changes that will occur as a result of reprovisioning the cluster.



Note: Any underprovisioned nodes are flagged in the Status column.

3. Click **Commit**.

The **Cluster Reprovisioned** dialog box opens, confirming that the disk pools have been provisioned.

4. Click **OK**.

Disk pool management

As nodes are added to an Isilon cluster, they are automatically provisioned into disk pools based on similar attributes such as node type or disk size. If automatic provisioning is not sufficient for your workflow, you can manually create and provision custom disk pools.



Note: All nodes in the cluster must be under an Isilon support contract; if any node is not under an Isilon support contract, none of the nodes in the cluster will be supported.

Monitor disk pools

Single-node and two-node disk pools are considered *underprovisioned* and are not supported. You can view the status of each disk pool in the cluster to determine whether additional provisioning is required.

1. On the **File System** menu, point to **SmartPools**, and then click **Disk Pools**.
The SmartPools **Disk Pools** page appears.
2. In the **Disk Pools** section, you can determine the status of a disk pool in the cluster by viewing or hovering over its icon in the **Status** column. A disk pool can be in one of the following states:
 - **Provisioned:** A green status icon indicates that the disk pool is properly provisioned.
 - **Underprovisioned:** A gray status icon indicates that the disk pool is empty or underprovisioned. You can fully provision the disk pool by adding or reallocating nodes.

Create a disk pool

This procedure explains how to create a disk pool and assign nodes to the pool.



Note:

- You cannot create an empty disk pool.
 - Each disk pool requires at least three nodes to be fully provisioned. An underprovisioned disk pool will not be populated until at least three nodes are assigned to it.
 - You cannot split the drives in a node across multiple disk pools.
1. On the **File System** menu, point to **SmartPools**, and then click **Disk Pools**.
The SmartPools **Disk Pools** page appears.
 2. In the **Disk Pools** section, click **Manually add disk pool**.
The SmartPools **Manually Add Disk Pool** page appears.
 3. In the **Disk Pool Attributes** section, set the following values:
 - **Pool name:** Specify a name for the disk pool. This setting is required.
 - **Protection level:** Optionally select the protection policy for the disk pool.
 4. In the **Manual Disk Pool Allocation** section, add or remove resources as needed. You must assign at least one node to successfully create the pool.
 - To add a node to the disk pool, under **Other Resources**, select the node and then click the left-arrow button.



Note: You can optionally filter the list of available nodes by specifying any combination of series, platform, or disk pool.

- To remove a node from the disk pool, under **Current Resources**, select the node and then click the right-arrow button.

5. Click **Submit**.

A confirmation dialog box appears.

6. In the **Confirm** dialog box, click **Yes**.



Note: If changes to the disk pool allocation result in underprovisioned disk pools, a second confirmation dialog box appears. Click **Yes** to continue anyway or **No** to cancel the operation and reassign resources.

Modify a disk pool

You can rename disk pools, modify their protection policies, and manually reallocate nodes from one disk pool to another.

1. On the **File System** menu, point to **SmartPools**, and then click **Disk Pools**.
The SmartPools **Disk Pools** page appears.
2. In the **Disk Pools** section, under **Actions**, click **Edit** for the disk pool that you want to modify.
The SmartPools **Edit Disk Pool** page appears.
3. Modify the disk pool's settings as needed.
4. Click **Submit**.
A confirmation dialog box appears.
5. In the **Confirm** dialog box, click **Yes**.



Note: If changes to the disk pool allocation result in one or more underprovisioned disk pools, a second confirmation dialog box appears. Click **Yes** to continue or **No** to cancel the operation and reassign resources.

Delete a disk pool

You can delete a disk pool if it is no longer needed.

1. On the **File System** menu, point to **SmartPools**, and then click **Disk Pools**.
The SmartPools **Disk Pools** page appears.
2. In the **Disk Pool Status** section, under **Actions**, click **Delete** for the disk pool that you want to delete.
The **Confirm Delete Disk Pool** dialog box appears.



Note: If the disk pool is not empty, the **Cannot Delete Disk Pool** dialog box appears. Modify the disk pool to remove all resources and then retry.

3. Click **Yes** to confirm the deletion.
The disk pool is deleted from the cluster and removed from the list.

File pool policy management

File pools enable you to easily aggregate large numbers of files. You can add, modify, prioritize, copy, or remove file pool policies; modify the default policy; and apply and modify template policies.

File pool policies are managed from the SmartPools **File Pool Policies** page.



Note: If the **SmartPools is not managing protection settings** or **SmartPools is not managing I/O optimization settings** message displays, these settings will have no effect unless you enable SmartPools to manage those attributes via the SmartPools **Settings** page.

View file pool policies

The SmartPools **File Pool Policies** page displays currently configured file pool policies, available template policies, and the latest SmartPools job results.

On the **File System** menu, point to **SmartPools**, and then click **File Pool Policies**.

The SmartPools **File Pool Policies** page appears, and displays the following sections:

- **File Pool Policies:** Lists all configured file pool policies in the order in which they are processed. Details include the policy name and optional description; filter criteria, operation, and file-matching status; and available actions. In this section you can add, modify, delete, and copy file pool policies.
- **Template Policies:** Lists all available templates that you can use to create new file pool policies. Details include the template name and description; filter criteria and operation; and a link to use the template.
- **Latest Job Results:** Displays detailed results from the last time the SmartPools job was run.

Add a file pool policy

You can configure file pool policies to filter files according to criteria that you specify, and set protection and I/O optimization settings for files that match those criteria.

You can define zero or more operations per policy.

1. On the **File System** menu, point to **SmartPools**, and then click **File Pool Policies**.

The SmartPools **File Pool Policies** page appears.

2. In the **File Pool Policies** section, click **Add file pool policy**.

The SmartPools **Add File Pool Policy** page appears.

3. In the **Basic Settings** section, configure the following settings:

- **Policy name:** Type a name for the policy. This setting is required.
- **Description:** Optionally, type a description of the policy.

4. In the **Filter Settings** section, in the **Filter** box, click **Add criteria**.

The **Configure File Matching Criteria** dialog box appears.

- a. Click to select a filter.
- b. Click to select an available operator.
Operators vary according to the selected filter.
- c. Configure the comparison value.
Values vary according to the selected filter and operator.
- d. Click **Add**.



Note: The file pool policy requires at least one criterion. You can add additional criteria (one at a time) or remove criteria by clicking one of the following links.

- **Add 'AND' criteria:** Adds a criterion to the selected criteria block. Files must satisfy each criterion to match the filter.
- **Add 'OR' criteria:** Adds a criterion to a new criteria block. Files can satisfy either criteria block to match the filter. You can configure up to three criteria blocks per file pool policy.
- **Delete:** Removes the selected criterion.
- **Delete criteria block:** Removes an entire block of criteria.

5. Optionally check or clear the **Stop processing more rules for files matching this filter** check box.

By default, this option is selected and files that match the specified filter settings will not be processed by any other rules.

6. In the **Protection Settings** section, optionally configure the protection settings for files that are processed by this file pool policy. Otherwise, these settings are defined by the default file pool policy.



Note: If the **SmartPools is not managing protection settings** message displays, these settings will have no effect unless you enable SmartPools to manage protection settings via the SmartPools **Settings** page.

- To specify the data and snapshot pools, click the **Set data pool to** check box, and then configure the following settings for the data and snapshot pools.
 1. Select **any pool** to assign filtered files to disk pools without restriction, or manually select a pool from the list. For the snapshot pool, you can select **same as data pool**.
 2. If the pool you selected contains SSDs, select one of the following SSD strategies:
 - **Metadata acceleration:** Creates a mirror backup of file metadata on an SSD and writes the rest of the metadata plus all user data on HDDs. Depending on the global namespace acceleration setting, the SSD mirror may be an extra mirror in addition to the number required to satisfy the protection level.
 - **Avoid SSDs:** Never uses SSDs; writes all associated file data and metadata to HDDs only.
 - **Data on SSDs:** Similar to metadata acceleration, but also writes one copy of the file's user data (if mirrored) or all of the data (if not mirrored) on SSDs. Regardless of whether global namespace acceleration is enabled, any SSD blocks reside on the file's target pool if there is room. This SSD strategy does not result in the creation of additional mirrors beyond the normal protection level.
- To specify the protection level, click the **Set protection level to** check box, and then select one of the following options:
 - **Default protection level of disk pool:** Assigns the default protection policy of the disk pool to the filtered files.
 - **Specific level:** Assigns a specified protection policy to the filtered files. To change the protection policy, select a new value from the list.

7. In the **I/O Optimization Settings** section, optionally configure the SmartCache and data access pattern settings for files that are processed by this policy. Otherwise, these settings are defined by the default file pool policy.



Note: If the **SmartPools is not managing I/O optimization settings** message displays, these settings will have no effect unless you enable SmartPools to manage I/O optimization settings via the SmartPools **Settings** page.

- To enable or disable SmartCache for files that are processed by this policy, click the **Set SmartCache to** check box, and then select one of the following options:
 - **Enabled:** Enables SmartCache.



Note:

SmartCache can accelerate the process of writing content to the cluster. However, writes to the cluster using SMB will not be affected by SmartCache when write through is specified by the client. Also, writes to the cluster using NFS will not be affected by SmartCache if they are tagged as stable.

If a node crashes with SmartCache enabled while using NFS, unstable writes that have not been committed will be temporarily lost on the node. However, the uncommitted data will be sent to the cluster again as soon as the client reconnects to the cluster, unless the client crashes. As long as the client does not crash before it reconnects, no data will be lost.

- **Disabled:** Disables SmartCache.
- To specify the data access pattern for files that are processed by this policy, click the **Set data access pattern to** check box, and then select one of the following options:

- **Random:** Specifies a random access pattern.
- **Concurrency:** Specifies a concurrent access pattern.
- **Streaming:** Specifies a streaming access pattern. Streaming access patterns can improve performance in some workflows.

8. Click **Submit**.

The file pool policy is added to the File Pool Policies list. Changes will be applied the next time the SmartPools job runs.

9. If there are more than two file pool policies in addition to the default policy, optionally change the order in which they are processed by selecting the **Order** option for a policy and then clicking **Move up** or **Move down**.



Note: You cannot change the order of the default file pool policy.

10. Optionally click **Start SmartPools Job** to immediately enforce SmartPools file policies.



Note: Starting the SmartPools job may adversely impact cluster performance. If you do not want to start the SmartPools job immediately, it will start automatically at the next scheduled run time.

Modify a file pool policy

You can modify a file pool policy's name and description, filter criteria, and the protection and I/O optimization settings that will be applied to files that match the filter criteria.

1. On the **File System** menu, point to **SmartPools**, and then click **File Pool Policies**.
The SmartPools **File Pool Policies** page appears.
2. In the **File Pool Policies** section, under **Actions**, click **Edit** for the file pool policy that you want to modify.
The SmartPools **Edit File Pool Policy** page appears.
3. Modify the file pool policy's settings as needed.
4. Click **Submit**.
Changes will be applied the next time the SmartPools job runs.
5. Optionally click **Start SmartPools Job** to immediately enforce SmartPools file policies.



Note: Starting the SmartPools job may adversely impact cluster performance. If you do not want to start the SmartPools job immediately, it will start automatically at the next scheduled run time.

Prioritize a file pool policy

File pool policies are evaluated based upon the order in which they appear in the File Pool Policies list. By default, new policies are added to the end of the list and are evaluated last. You can give a policy higher or lower priority by moving it up or down the list.

1. On the **File System** menu, point to **SmartPools**, and then click **File Pool Policies**.
The SmartPools **File Pool Policies** page appears.
2. Repeat the following substeps until the policies are listed in the order that you want them evaluated:
 - a. In the **File Pool Policies** section, under **Order**, select the radio button for the policy whose order you want to change.
 - b. Click **Move up** or **Move down** until the policy is in the order that you want it.

Copy a file pool policy

You can create an exact copy of any file pool policy except the default policy, and then modify the policy's settings if required.

1. On the **File System** menu, point to **SmartPools**, and then click **File Pool Policies**.
The SmartPools **File Pool Policies** page appears.
2. In the **File Pool Policies** section, under **Actions**, click **Copy** for the file pool policy that you want to copy.
A copy of the policy is created with the name "Copy of <Policy Name>."
3. Optionally modify the policy's settings as needed.
For more information, see "Modify a file pool policy."

Remove a file pool policy

You can remove any file pool policy except the default policy.

1. On the **File System** menu, point to **SmartPools**, and then click **File Pool Policies**.
The SmartPools **File Pool Policies** page appears.
2. In the **File Pool Policies** section, under **Actions**, click **Delete** for the file pool policy that you want to remove.
A confirmation dialog box appears.
3. Click **Yes**.
The file pool policy is removed from the list.

Use a file pool policy template

SmartPools includes templates for file pool policies, which you can use as-is or as a starting point for further configuration.

You cannot modify, remove, or create new file pool policy templates.

1. On the **File System** menu, point to **SmartPools**, and then click **File Pool Policies**.
The SmartPools **File Pool Policies** page appears.
2. In the **Template Policies** section, under **Actions**, click **Use** for the template policy that you want to use.
The SmartPools **Use File Pool Policy Template** page appears.
3. Optionally modify the policy's settings as needed.
For more information, see "Modify a file pool policy."
4. Click **Submit**.
The SmartPools **File Pool Policies** page appears, and the policy is added to the **File Pool Policies** list.

Modify the default file pool policy

The default file pool policy is defined by the global SmartPools settings. You can modify the default policy by configuring the settings on the SmartPools **Settings** page.

This procedure describes how to access the SmartPools **Settings** page from the SmartPools **File Pool Policies** page. You can also access this page by clicking the **File System** menu, pointing to **SmartPools**, and then clicking **Settings**.



Note: The default file pool policy is always last in the ordered list of enabled policies. You cannot reorder or remove the default file pool policy.

1. On the **File System** menu, point to **SmartPools**, and then click **File Pool Policies**.
The SmartPools **File Pool Policies** page appears.
2. In the **File Pool Policies** section, under **Actions**, click **Modify defaults** for the default file pool policy.
The SmartPools **Settings** page appears.
3. Modify the settings as needed.

For more information, see "SmartPools configuration."

Software module licensing

The Isilon OneFS file system includes selected software applications, or modules, with the base system. You can optionally purchase and enable additional, separately licensed software modules that provide enhanced features. Each optional Isilon software module requires a separate license.

To activate an Isilon module, you must obtain a valid license key and then enter the key through the Isilon web administration interface or the command-line interface.

For more information about optional software modules, or to obtain a module license, contact your Isilon sales representative.

View module license status information

You can view information about the current status of any optional Isilon software modules.

1. On the **Help** menu, click **Versions and Licenses**.
The **Versions and Licenses** page appears.
2. In the **Licensed Modules** area, review the status information about each optional software module:
 - **Module** column: Displays the name of the software module.
 - **Status** column: Indicates the current license state of the module.
 - **Configuration** column: Displays configuration information about the module.
 - **Expires** column: Indicates when the module license is scheduled to expire, or when it expired.

Activate a module license through the web interface

To activate an optional Isilon software module, you must enter a valid module license key through the Isilon web administration interface or the command-line interface. This section describes how to activate a module license through the web administration interface.

Prerequisite: Before you can activate an Isilon software module, you must obtain a valid license key. For information about requesting and receiving license keys, contact your Isilon sales representative.

1. On the **Cluster** menu, point to **Cluster Management**, and then click **Activate License**.
The **Activate License** page appears.
2. In the **License key** box, type the license key for the module that you want to enable.
3. Read the end user license agreement, click **I have read and agree...**, and then click **Submit**.
The **License List** page appears, displaying updated status information about the software module.

Activate a module license through the command-line interface

To activate an optional Isilon software module, you must enter a valid module license key through the Isilon web administration interface or the command-line interface. This section describes how to activate a module license through the command-line interface.

Prerequisite: Before you can activate an Isilon software module, you must obtain a valid license key. For information about requesting and receiving license keys, contact your Isilon sales representative.

1. Open a secure shell (SSH) connection with any node in the cluster, and then log in to the cluster using the "root" account.
2. At the command prompt, enter the following command, replacing **<license key>** with the appropriate license key for the module.

```
isi license activate <license key>
```

The system returns a confirmation message similar to the following output:

```
SUCCESS,  
<module name> has been successfully activated.
```

Unconfigure a module license

You may want to unconfigure a license for an Isilon software module if, for example, you enabled an evaluation version of a module but later decided not to purchase a permanent license.

Unconfiguring a module license disables any recurring jobs or scheduled operations that you have enabled for that module. The specific results of unconfiguring a license are different for each module:

- **iSCSI:** If you unconfigure an iSCSI license, iSCSI initiators can no longer establish iSCSI connections to the cluster.
- **SmartConnect:** If you unconfigure a SmartConnect license, the system converts any dynamic IP address pools to static IP address pools.
- **SmartPools:** If you unconfigure a SmartPools license, the system disables all configured file pool policies except the default file pool policy. The SmartPools job is replaced by the SetProtectPlus job, which runs only when needed and applies to only the default file pool policy.
- **SmartQuotas:** If you unconfigure a SmartQuotas license, the system disables any configured thresholds.
- **SnapshotIQ:** If you unconfigure a SnapshotIQ license, the system disables any configured snapshot schedules.
- **SyncIQ:** If you unconfigure a SyncIQ license, the system disables any configured SyncIQ policies and jobs.

You can unconfigure module licenses only through the command-line interface. You cannot unconfigure a module license through the web administration interface.

1. Open a secure shell (SSH) connection to any node in the cluster and log in using the root user account.
2. At the command prompt, run the following command, where **<module_name>** is the name of the module to unconfigure.

```
isi license unconfigure <module_name>
```

The supported **<module_name>** values are **iscsi**, **smartconnect**, **smartpools**, **smartquotas**, **snapshotiq**, or **synciq**.

The Isilon SyncIQ module

The Isilon SyncIQ software module is an optional tool that enables you to flexibly manage and automate data synchronization between two Isilon IQ clusters. The SyncIQ module can copy data from one Isilon cluster to another, or synchronize data by replicating files from a source cluster to a target cluster and then deleting any files on the target that are not present on the source.

You can configure one or more SyncIQ policies, which determine how and when SyncIQ jobs will run. When SyncIQ transfers updates to a file that was transferred during a previous SyncIQ job, SyncIQ transfers only the blocks of that file that have changed since the last job, rather than transferring the entire file.

The SyncIQ module requires a separate license. For information about enabling SyncIQ on your cluster, contact your Isilon sales representative.

SyncIQ snapshot overview

To provide point-in-time data protection, the Isilon SyncIQ module automatically generates a snapshot of the data set on the source cluster when a SyncIQ job starts. Optionally, if the Isilon SnapshotIQ module is licensed and enabled on the target cluster, you can configure SyncIQ to also take a snapshot of the data set on the target cluster after a SyncIQ job runs.

Source-cluster snapshots

The SyncIQ module automatically takes a snapshot of the data set on the source cluster before starting each SyncIQ data-synchronization or copy job; this source-cluster snapshot does not require a SnapshotIQ module license.

When a SyncIQ job starts, if the system detects a previous source-cluster snapshot, SyncIQ sends to the target only any files that are not present in that previous snapshot, as well as any files that have changed since the last source-cluster snapshot was taken.

When a SyncIQ job starts, if the system does not find a previous source-cluster snapshot (for example, if a SyncIQ job is running for the first time), SyncIQ takes an initial full snapshot of the specified root path on the source cluster.

When a SyncIQ job completes, the system deletes the previous source-cluster snapshot and retains the most recent snapshot until the next time the job runs.

Source-cluster snapshots are named `SIQ-<policy-id>-[new,latest]`, where *<policy-id>* is the unique system-generated policy identifier. Each time a SyncIQ job finishes running, the existing `latest` snapshot is replaced with the most recent snapshot, which in turn becomes the `latest` snapshot. The `new` label is temporary, and is used only while the job is running.

During a SyncIQ job, SyncIQ identifies any changes on the source cluster and then replicates those changes to the target cluster. SyncIQ does not support bi-directional synchronization and cannot detect files on the target directory. To prevent inadvertent data loss or modifications on the target, only SyncIQ can write to the target directory.

Target-cluster snapshots

If you want to generate a snapshot of the target cluster after each SyncIQ job finishes running, you must license and enable the Isilon SnapshotIQ module on the target cluster. Snapshots taken on the target can provide a consistent view of the synchronized files even during subsequent transfers, in addition to the data-retention and backup benefits that regular source-cluster snapshots provide.

To obtain a SnapshotIQ license for your target cluster, contact your Isilon sales representative.

SyncIQ policies and jobs

SyncIQ policies determine how and when SyncIQ jobs will run. Each SyncIQ policy defines the settings for one job profile; however, the job itself might run multiple times, depending on how you configure the policy's schedule settings or how often you choose to manually run the job.

You can create an unlimited number of SyncIQ policies, depending on your data-synchronization needs. However, you must create at least one policy before SyncIQ can copy or synchronize data.

You can configure detailed SyncIQ policies that define the source-cluster directories, the target-cluster path, the policy action (copy or synchronize), file criteria, target-cluster snapshot settings, and other, advanced settings. Each policy runs associated copy or synchronize jobs according to a schedule that you specify during the policy configuration process. Alternatively, you can configure unscheduled jobs, which run only if you manually start them.

A maximum of five SyncIQ jobs can run concurrently on a source cluster. You can start more than five jobs, or schedule more than five jobs to run during the same time period. However, only the first five jobs will run immediately; the remaining jobs wait in a queue until it is their turn to start. Note, however, that an unlimited number of jobs can run concurrently against a target cluster.

Each SyncIQ policy has a direct association with its specified target directory, and only SyncIQ can write to the specified path. This prevents the inadvertent modification, creation, or deletion of files in the policy's specified target.

Source and target cluster association

OneFS associates a policy with its specified target directory by placing a cookie on the source cluster when the job runs for the first time. Even if you modify the name or IP address of the target cluster, the cookie causes the association to persist. If necessary, you can manually break a target association, for example if an association is obsolete or was intended for temporary testing purposes. Breaking a target association causes the source to fully resynchronize the next time the job runs; during this full resynchronization, SyncIQ creates a new association between the source and its specified target.



Caution: Depending on the amount of data being synchronized or copied, a full resynchronization can take a very long time to complete.

If the association between the source and target clusters is broken, the policy falls into an unrunnable and disabled state. If a policy falls into an unrunnable state, SyncIQ prompts you to resolve the association between the source and the target before the job can run again.

In some cases, you may want to preemptively and manually reset the association between a source and a target, for example if you have modified an existing policy to point at a different target. You can reset an association for a healthy policy at any time.



Note: Resetting a policy causes a full synchronization to occur.

You can break an association from the target cluster. However, you can reset or resolve an association from only the source cluster.

SyncIQ policy configuration

SyncIQ policies define how and when data will be copied or synchronized. You can create one or more SyncIQ policies, depending on your data-synchronization needs. You must create at least one policy in order to use the SyncIQ module.

Configuring a SyncIQ policy is a seven-step process:

1. Configure basic SyncIQ policy settings.
2. Configure source-cluster settings.
3. Configure file-criteria settings.
4. Configure basic target-cluster settings.
5. Optionally, configure target-cluster snapshot settings.

6. Optionally, configure advanced settings.
7. Review and then save the policy settings.



Important: When creating a new SyncIQ policy, plan and configure the initial settings carefully. If you later modify any of the following settings in an existing SyncIQ policy, the system will perform a *full* resynchronization of all affected data the next time the associated job runs:

- Any source-cluster directories, including the root path and any explicitly included or excluded directories
- Target path
- File criteria

Modifying any other SyncIQ policy settings, including most target-cluster settings, does not cause a full resynchronization to occur. However, if you specify a different physical target cluster or a new target path, a full resynchronization will occur, even if the new cluster's name is identical to the previous cluster's name.

Excluding or including source-cluster directories

Each time a SyncIQ policy runs a job, the system takes a snapshot of the specified root directory. This snapshot becomes the basis for the replication operation between the two specified clusters.

When configuring source-cluster settings in a SyncIQ policy, in addition to specifying a root directory on the source cluster, you can optionally exclude or include specific source-cluster directories. By default, all files and folders under the specified root directory are synchronized to the target cluster during a SyncIQ job. However, if you explicitly exclude any directories, those directories, and any files contained in them, are not synchronized to the target cluster. In addition, if you explicitly include any directories in the policy configuration, the system synchronizes only the files that are contained in that explicitly included directory to the target cluster.

Any directories that you explicitly include or exclude must be contained in or under the specified root directory. For example, consider a policy in which the specified root directory is `/ifs/data`. In this example, you could explicitly include the `/ifs/data/media` directory because it is under `/ifs/data`. When the associated policy runs, only the contents of the `/ifs/data/media` directory would be synchronized to the target cluster.

If you explicitly exclude a directory that is contained in the specified root directory, and you do not explicitly include any directories, the contents of the excluded directory are not synchronized to the target cluster. If, however, you explicitly include directories, any explicitly excluded directories must be contained in one of the explicitly included directories; otherwise, the excluded-directory setting has no effect. For example, consider a policy in which the specified root directory is `/ifs/data`, and the following directories are explicitly included and excluded:

Explicitly included directories:

- `/ifs/data/media/music`
- `/ifs/data/media/movies`

Explicitly excluded directories:

- `/ifs/data/media/music/working`
- `/ifs/data/media`

In this example, the setting that explicitly excludes the `/ifs/data/media` directory has no effect because the `/ifs/data/media` directory is not contained under either of the explicitly included directories.

In addition, if you exclude a directory that contains the specified root directory, that exclude-directory setting has no effect. For example, consider a policy in which the specified root directory is `/ifs/data`. Configuring a policy setting that excludes the `/ifs` directory has no effect, and all of contents of the specified root directory (in this example, `/ifs/data`) are synchronized to the target cluster.

Configure basic SyncIQ policy settings

Configuring basic policy settings is the first of seven steps required to configure a SyncIQ policy.

1. On the **File System** menu, point to **SyncIQ**, and then click **Policies**.
The SyncIQ **Policies** page appears.

2. Click **Add policy**.

The SyncIQ **Add Policy** page appears.



Note: Depending on whether you are creating a new SyncIQ policy or modifying an existing policy, the name of this page is either **SyncIQ Add Policy** or **SyncIQ Edit Policy**.

3. In the **Basic Settings** section, in the **Policy name** box, type a name for the policy.
Policy names can include only letters, numerals, hyphens, and underscore characters.
4. In the **Description** box, type a descriptive comment about the policy.
This step is optional, but a description can be helpful later when reviewing and managing your policies.
5. In the **Action** area, specify a job type:
 - To copy data from a source cluster to a target cluster, click **Copy**.
 - To synchronize data between two clusters, click **Synchronize**. If you select this option, data on the specified source cluster is replicated to the specified target cluster when the job runs, and any files on the target that are not present on the source are deleted.
6. In the **Run job** area, specify whether the SyncIQ job will run manually or automatically:
 - **Manually:** Click to configure the SyncIQ job to run only when you manually initiate it. You can manually run a SyncIQ job by clicking the policy's **Start** link on the **SyncIQ Policies** page.
 - **Scheduled:** Click to configure the SyncIQ job to run automatically according to a schedule, and then click **Edit schedule** to configure the **Interval** and **Frequency** settings.



Note: Scheduled SyncIQ policies with a **Frequency** setting of **Multiple times** cannot span across more than one calendar day. For example, a SyncIQ policy that runs every five minutes starting at 7:00 PM and ending at 11:00 PM is valid; however, a policy that runs every five minutes starting at 7:00 PM and ending at 1:00 AM on the following day is invalid.

To save the policy, you must complete each of the remaining steps in the SyncIQ policy-configuration process. The next step in the process is configuring source-cluster settings. For more information, see [Configure SyncIQ policy source-cluster settings](#) on page 142.

Configure SyncIQ policy source-cluster settings

For each SyncIQ policy, you must specify one or more source directories on the source cluster.

Prerequisite: This is the second of seven steps required to configure a SyncIQ policy. Before you configure source-cluster settings, you must configure basic SyncIQ policy settings. For more information, see [Configure basic SyncIQ policy settings](#) on page 141.

1. On the SyncIQ **Add Policy** page, in the **Source Cluster** area, type the full path (beginning with `/ifs`) of the source directory in the **Root directory** box, or click **Browse** to locate it.

The root directory must be contained in the `/ifs/` directory. You cannot specify any path that ends with `.snapshot` as the root directory.



Note: Depending on whether you are creating a new SyncIQ policy or modifying an existing policy, the name of this page is either **SyncIQ Add Policy** or **SyncIQ Edit Policy**.

2. In addition to the specified root directory, which is included by default, you can optionally include or exclude specific directories that are contained in the root directory by clicking the **Add directory** links to the right of the **Include directories** and **Exclude directories** boxes.
3. Optionally, repeat steps 1 and 2 as needed.

To save the policy, you must complete each of the remaining steps in the SyncIQ policy-configuration process. The next step in the process is configuring file-criteria settings. For more information, see [Configure SyncIQ policy file criteria](#) on page 143.

Configure SyncIQ policy file criteria

For each SyncIQ policy, you can define file-criteria statements that explicitly include or exclude files from the policy action.

Prerequisite: This is the third of seven steps required to configure a SyncIQ policy. Before you configure file criteria, you must configure source-cluster settings. For more information, see [Configure SyncIQ policy source-cluster settings](#) on page 142.

A file-criteria statement can include one or more elements. Each file-criteria element contains a file attribute, a comparison operator, and a comparison value. To combine multiple criteria elements into a criteria statement, use the Boolean *AND* and *OR* operators. You can configure any number of *AND* and *OR* file-criteria definitions. If you do not specify any file-criteria statements, by default all files in the source cluster's specified paths are included in the policy action.



Note: Configuring file-criteria statements can cause the associated jobs to run slowly. It is recommended that you specify file-criteria statements in a SyncIQ policy only if necessary.



Important: For synchronization policies, use caution if you modify a file attribute's comparison operators and comparison values. Modifying these settings will cause any non-matching files to be deleted from the target the next time the job runs. This does not apply to policies that copy data.

1. On the SyncIQ **Add Policy** page, in the **Source Cluster** area, review the default **File criteria** settings. The default file-criteria definition is **File name is equal to '*'**, where the wildcard character * represents any value.



Note: Depending on whether you are creating a new SyncIQ policy or modifying an existing policy, the name of this page is either **SyncIQ Add Policy** or **SyncIQ Edit Policy**.

If necessary, you can remove a file-criteria definition by clicking the **Delete** link to its right, or you can delete an entire block of *AND* or *OR* file-criteria definitions by clicking **Delete criteria block**.

2. If you want to configure more specific file-criteria definitions, click **Add criteria** or **Add 'AND' criteria**, and then proceed to step 3. (If you do not want to configure additional file-criteria definitions, do not complete the rest of this procedure.)

The **Configure File Matching Criteria** dialog box appears.

3. In the list on the left, click the appropriate file attribute:
 - **Date created:** Specifies file criteria based on when the file was created. This option is available for only SyncIQ copy policies; it does not apply to synchronization policies.
 - **Date accessed:** Specifies file criteria based on when the file was last accessed. This option is available for only SyncIQ copy policies; it does not apply to synchronization policies. This setting is available only if the cluster's global access-time-tracking option is enabled.
 - **Date modified:** Specifies file criteria based on when the file was last modified in increments of days, weeks, months, or years. This option is available for only SyncIQ copy policies; it does not apply to synchronization policies.
 - **File name:** Specifies file criteria based on the file name.
 - **Path:** Specifies file criteria based on where the file is stored. This option is available for only SyncIQ copy policies; it does not apply to synchronization policies.
 - **Size:** Specifies file criteria based on the file size.
 - **Type:** Specifies file criteria based on the file-system object type (**Regular file**, **Directory**, or **Soft link**).

The supported comparison-operator options for the selected file attribute appear in the comparison-operator list, to the right of the file-attribute list.

4. Click the appropriate comparison operator and comparison value:

- **Date created** options: Specify a **Relative date** integer value (in days, weeks, months, or years ago) or a **Fixed date** value that references a specific date and time.
- **Date accessed** options: Specify a **Relative date** integer value (in days, weeks, months, or years ago) or a **Fixed date** value that references a specific date and time. Time settings are based on a 24-hour clock.
- **Date modified** options: Specify a **Relative date** integer value (in days, weeks, months, or years ago) or a **Fixed date** value that references a specific date and time.
- **File name** options: Click **is** (to include the specified text) or **is not** (to exclude the specified text), and then type a full or partial file name in the **Basic** box. You can include the wildcard characters *****, **?**, and **[]**.



Note: Alternatively, if you want to specify more detailed information, you can click **Advanced** and then type POSIX regular-expression (regex) text in the **Advanced** box. Regular expressions are sets of symbols and syntactic elements that are used to match patterns of text. These expressions can be more powerful and flexible than simple wildcard characters. Isilon clusters support IEEE Std 1003.2 (POSIX.2) regular expressions. For example, to select all files ending in `.jpg`, you could type `.*\.jpg$`. To select all files with either `.jpg` or `.gif` file extensions, you could type `.*\.(jpg|gif)$`. For more information about POSIX regular expressions, see the BSD man pages.

- **Path** options: Click **is** (to include the specified path) or **is not** (to exclude the specified path), and then type a full or partial path. You can include the standard wildcard characters *****, **?**, and **[]** in the path box.
- **Size** options: Specify the appropriate comparison operator, and then type an integer that represents the file size in bytes, **KB**, **MB**, **GB**, or **TB**.



Note: File sizes are represented in multiples of 1024, not 1000.

- **Type** options: Click **is** (to include the specified file-system object type) or **is not** (to exclude the specified object type), and then click the type of object to include or exclude. The supported file-system object types are **Regular file**, **Directory**, or **Soft link**.



Note: A soft link is a special type of POSIX-supported file that contains a reference to another file or directory.

5. Click **Done**.

The new file-criteria statement appears in the **File criteria** list.

6. Optionally, repeat steps 2 through 5 as needed to add additional **AND** elements.

7. Optionally, create **OR** elements as needed:

a. Click **Add 'OR' criteria**.

The **Configure File Matching Criteria** dialog box appears.

b. Repeat steps 3 through 5 as needed to add one or more **OR** elements.

8. Optionally, repeat steps 2 through 7 as needed.

To save the policy, you must complete each of the remaining steps in the SyncIQ policy-configuration process. The next step in the process is configuring basic target-cluster settings. For more information, see [Configure SyncIQ policy basic target-cluster settings](#) on page 144.

Configure SyncIQ policy basic target-cluster settings

For each SyncIQ policy, you must specify a target directory on the target cluster.

Prerequisite: This is the fourth of seven steps required to configure a SyncIQ policy. Before you configure target-cluster settings, you must configure file-criteria settings. For more information, see [Configure SyncIQ policy file criteria](#) on page 143.

1. On the SyncIQ **Add Policy** page, in the **Target Cluster** area, type one of the following values in the **Name or address** box:

- the host name of any node in the target cluster

- the fully qualified domain name of any node in the target cluster
- the name of a SmartConnect zone in the target cluster
- the IPv4 or IPv6 address of any node in the target cluster

Type the **Name or address** value carefully; the system does not validate the data that you enter.



Note: SyncIQ does not support dynamic pools.



Note: Depending on whether you are creating a new SyncIQ policy or modifying an existing policy, the name of this page is either SyncIQ **Add Policy** or SyncIQ **Edit Policy**.

2. Click one of the following options to specify which nodes on the target cluster to connect to when the SyncIQ job runs:

These settings are applicable only if you specified a host name, domain name, or SmartConnect zone name in the target cluster **Name or address** box. They are not applicable if you specified an IP address.

- **Connect to any nodes in the cluster:** Connect to any available nodes on the target cluster.
- **Connect to only the nodes in the subnet and pool if the target cluster name specifies a SmartConnect zone:** Connect to only the target-cluster nodes that are contained in the subnet and pool that you specified in the target cluster **Name or address** box, if you specified a SmartConnect zone.

3. In the **Target directory** box, type the absolute path (beginning with `/ifs`) for the target-cluster directory to which files will be replicated.



Caution: Do not specify a target directory that is already specified as the target directory of a different SyncIQ policy.



Note: In most cases, you should not specify `/ifs` as the target directory; doing so would prevent you from using the target cluster for any purposes other than as a SyncIQ target because the entire cluster would be put into a read-only state and all of the data contained in `/ifs` would be overwritten each time the job runs.

If the specified target directory does not already exist on the target cluster, the directory will be created the first time the job runs. For synchronization jobs, if the directory already exists on the target cluster, all files in that directory will be deleted. For copy jobs, target files may be overwritten with source files.

To save the policy, you must complete each of the remaining steps in the SyncIQ policy-configuration process. The next step in the process is configuring target-cluster snapshot settings. For more information, see [Configure SyncIQ policy target-cluster snapshot settings](#) on page 145.

Configure SyncIQ policy target-cluster snapshot settings

For each SyncIQ policy, you can optionally configure SyncIQ to save a snapshot of the data set on the target cluster after each job is complete.

Prerequisite: This is the fifth of seven steps required to configure a SyncIQ policy. Before you configure target-cluster snapshot settings, you must configure general target-cluster settings. For more information, see [Configure SyncIQ policy basic target-cluster settings](#) on page 144. To use the target-cluster snapshot feature, the Isilon SnapshotIQ module must be licensed and enabled on the target cluster. To obtain a SnapshotIQ license for your target cluster, contact your Isilon sales representative.

1. On the SyncIQ **Add Policy** page, in the **Target Cluster Snapshot** area, configure target-cluster snapshot settings as necessary.



Note: Depending on whether you are creating a new SyncIQ policy or modifying an existing policy, the name of this page is either SyncIQ **Add Policy** or SyncIQ **Edit Policy**.

2. If you want to automatically create a snapshot on the target cluster after each job completes, in the **Snapshots** section, click **Create a snapshot on the target cluster**.



Note: Isilon recommends enabling the **Create a snapshot on the target cluster** option. If this option were disabled and the source cluster were to fail during a synchronization job, data might be only partially written to the target. This could leave the clusters in an inconsistent state and could result in corrupted files. Enabling the **Create a snapshot on the target cluster** option ensures a consistent snapshot of a known good state.



Note: If you enable the **Create a snapshot on the target cluster** option, SyncIQ takes a snapshot of target-cluster data set after each job completes. In addition, if a previous target snapshot is not detected (for example, if a SyncIQ job is running for the first time), SyncIQ also takes a snapshot of the data set before the job runs.

3. In the **Snapshot alias name** box, type an alias-name string using the supported string variables that are listed to the right of the box.

An alias is an alternative name that is usually shorter or easier to enter than the schedule name. The default snapshot alias-name string is **SIQ-%<SrcCluster>-%<PolicyName>-latest**.

4. In the **Snapshot pattern** box, type a naming-pattern string for snapshots, using the supported string variables that are listed to the right of the box.

A snapshot pattern is a template that uses date variables to name generated snapshots with what is essentially a time-and-date stamp. The default naming-pattern string is

SIQ-%<SrcCluster>-%<PolicyName>-%Y-%m-%d_%H-%M.

5. In the **Snapshot expiration** section, specify whether and when to automatically delete snapshots from the target cluster after a specified interval:

- If you do not want snapshots to ever be automatically deleted, click **Do not delete snapshots**. You can manually delete or otherwise manage these snapshots on the target cluster.
- If you want snapshots to be deleted after a specified amount of time has elapsed, click **Delete snapshots when they expire** and then configure an **Expiration** schedule by typing an integer that represents (in **days**, **weeks**, **months**, or **years**) how long snapshots will be stored on the cluster before they are automatically deleted.

To save the policy, you must complete each of the remaining steps in the SyncIQ policy-configuration process. The next step in the process is configuring advanced SyncIQ policy settings. For more information, see [Configure advanced SyncIQ policy settings](#) on page 146.

Configure advanced SyncIQ policy settings

For each SyncIQ policy, you can define optional advanced settings. The default values for the advanced settings are sufficient for most SyncIQ deployments.

Prerequisite: This is the sixth of seven steps required to configure a SyncIQ policy. Before you configure advanced SyncIQ policy settings, you must configure target-cluster snapshot settings. For more information, see [Configure SyncIQ policy target-cluster snapshot settings](#) on page 145.



Note: The settings that appear in the **Advanced Settings** area of the SyncIQ **Add Policy** or **Edit Policy** page vary depending on the type of job action (copy or synchronize) you specified for the policy.

1. On the SyncIQ **Add Policy** page, in the **Advanced Settings** area, review the default advanced settings and then, if necessary, continue with this procedure to modify settings.



Note: Depending on whether you are creating a new SyncIQ policy or modifying an existing policy, the name of this page is either SyncIQ **Add Policy** or SyncIQ **Edit Policy**.

For most SyncIQ deployments, the default advanced settings are sufficient.

2. Review the **Workers per node** setting, which specifies the maximum number of concurrent processes per node that can be used to perform SyncIQ operations. The default setting is **3**.



Caution: Do not modify this default setting without assistance from Isilon Technical Support.

3. In the **Log level** list, click to specify the level of logging to perform for the operation:

- **Notice:** Logs job-level and process-level activity, including job starts and stops, and worker coordination information. This is the default log level, and is recommended for most SyncIQ deployments.
- **Error:** Logs only events related to specific types of failures.
- **Network Activity:** Logs expanded job-activity and work-item information, including specific paths and snapshot names.
- **File Activity:** Logs a separate event for each action taken on a file. Do not enable this logging level without assistance from Isilon Technical Support.

These SyncIQ logs are typically used only for debugging purposes. If necessary, you can log on to the appropriate node through the command-line interface and view the contents of the node's `/var/log/isi_migrate.log` file.

4. If you want OneFS to perform a checksum on each file data packet that is affected by the SyncIQ job, select the **Validate file integrity** check box. This setting is enabled by default.

If the checksum values do not match, SyncIQ retransmits the affected file data packet.

5. Optionally, in the **Shared secret** box, type a shared secret for SyncIQ operations.

This shared secret provides a simple level of authentication that can prevent certain types of attacks; however, this feature does not perform any encryption. To establish this type of authentication, you must configure both the target cluster and the source cluster to require the same shared secret. For more information, see the Isilon Knowledgebase.

6. In the **Keep reports for** box, type an integer that represents (in **days**, **weeks**, **months**, or **years**) how long reports will be stored on the cluster before they are automatically deleted.

This setting takes precedence over any global default report-retention setting you might have configured on the SyncIQ **Settings** page. By default, the cluster retains a maximum of 2000 SyncIQ reports per policy.



Note: Some units of time specified for report retention are displayed differently in the web administration interface than how you originally enter them. Entering a number of days that is equal to a corresponding value in weeks, months, or years results in the larger unit of time being displayed. For example, if you enter a value of **30 days**, the web interface displays that value as **1 month**, while **364 days** is represented as **52 weeks**, and **365 days** is represented as **1 year**. This change occurs because OneFS internally records report retention times in seconds and then converts them into days, weeks, months, or years for display.

7. Specify a **Source node restrictions** setting:

This setting takes precedence over any global default source node restriction setting you might have configured on the SyncIQ **Settings** page.

- **Run the policy on all nodes in this cluster:** Click to connect to all nodes on the source cluster when the job runs.
- **Run the policy only on nodes in the specified subnet and pool:** Click to connect to only the source-cluster nodes that are contained in a specified subnet and pool, and then click to select the subnet and pool that you want to enable connections to. The **Subnet and pool** list includes all configured subnet-and-pool combinations that have been assigned static IP addresses.



Note: SyncIQ does not support dynamic pools.

8. If you want to capture information about files that are deleted during synchronization jobs, in the **Delete on synchronization** section, click **Record when a synchronization deletes files or directories**.

This setting applies to only synchronization jobs; it does not apply to copy jobs.

To save the policy, you must complete the remaining step in the SyncIQ policy-configuration process. The next step in the process is reviewing and saving the policy settings. For more information, see [Save SyncIQ policy settings](#) on page 148.

Save SyncIQ policy settings

Reviewing and saving policy settings is the seventh and final step required to configure a SyncIQ policy.

Prerequisite: Before you can save your SyncIQ policy settings, you must configure advanced SyncIQ policy settings. For more information, see [Configure advanced SyncIQ policy settings](#) on page 146.

1. On the SyncIQ **Add Policy** page, review the policy's current settings.



Note: Depending on whether you are creating a new SyncIQ policy or modifying an existing policy, the name of this page is either **SyncIQ Add Policy** or **SyncIQ Edit Policy**.

2. If you need to make additional changes, modify policy settings as needed.
3. Click **Submit**.

The SyncIQ **Policies** page appears, and the policy appears in the policy list.

Configure global default SyncIQ policy settings

You can configure global policy settings that apply to any new SyncIQ policies by default. These global default settings are automatically applied to any new SyncIQ policies that you create. However, any settings that you configure for specific policies will take precedence over the global default policy settings.

Modifying the global default SyncIQ policy settings does not affect existing SyncIQ policies. These global default settings apply to only any new SyncIQ policies that you create after configuring these global settings.

1. On the **File System** menu, point to **SyncIQ**, and then click **Settings**.

The SyncIQ **Settings** page appears.

2. In the **Default Policy Settings** area, specify a **Target cluster restrictions** setting:

- **Connect to any nodes in the cluster:** Click to connect to all nodes on the target cluster when the job runs.
- **Connect to only the nodes in the subnet and pool if the target cluster name specifies a SmartConnect zone:** Click to connect to only the target-cluster nodes that are contained in the SmartConnect zone that you specified in the target cluster **Name or address** box for each policy, if you specified a SmartConnect zone.



Note: SyncIQ does not support dynamic pools.

3. Specify a **Source cluster restrictions** setting:

- **Run the policy on all nodes in this cluster:** Click to connect to all nodes on the source cluster when the job runs.
- **Run the policy only on nodes in the specified subnet and pool:** Click to connect to only the source-cluster nodes that are contained in a specified subnet and pool, and then click to select the subnet and pool that you want to enable connections to. The **Subnet and pool** list includes all configured subnet-and-pool combinations that have been assigned static IP addresses.



Note: SyncIQ does not support dynamic pools.

4. In the **Report Settings** area, in the **Retain reports for** box, type an integer that represents (in **days, weeks, months, or years**) how long reports will be stored on the cluster before they are automatically deleted.
5. Click **Submit**.

SyncIQ policy management

SyncIQ policy-management tasks include modifying, enabling, disabling, and deleting policies.

View SyncIQ policies

You can view a list of all SyncIQ policies, as well as individual settings for each policy. You can also view information about all currently running SyncIQ jobs and historical information about recently completed SyncIQ jobs.

1. On the **File System** menu, point to **SyncIQ**, and then click **Policies**.
The SyncIQ **Policies** page appears, and displays a list of all configured SyncIQ policies.
2. Review the policy data:
 - **Run:** Indicates the status (running or paused) of a currently running job, or indicates that SyncIQ is disabled.
 - **Data:** Indicates the status of the last run of the job. A green icon indicates that the last job completed successfully. A yellow icon indicates that the last job did not complete successfully, but that an earlier job did complete successfully. If no icon appears, a job for the policy has not run.
 - **Policy:** Displays the name of the policy. You can view or edit the policy's settings by clicking the policy name.
 - **Last Known Good:** Indicates when the last successfully completed job ran, if applicable.
 - **Schedule:** Indicates when the job is scheduled to run. A value of **Manual** indicates that the job can be run only manually.
 - **Source:** Displays the source path or paths.
 - **Target:** Displays the target path.
 - **Actions:** Displays any policy-related actions that you can perform. You can perform an action by clicking its associated link.

Enable or disable a SyncIQ policy

You can selectively enable or disable SyncIQ policies to accommodate your data-replication needs. As an alternative to permanently deleting a policy, you can instead disable the policy, which keeps it available for future use.

If you disable a policy, any pending jobs scheduled for that policy will not run.

Before you can manually run a SyncIQ job, you must enable the associated SyncIQ policy.



Note: You cannot enable or disable a SyncIQ policy if its associated job is currently running.

1. On the **File System** menu, point to **SyncIQ**, and then click **Policies**.
The SyncIQ **Policies** page appears, and displays a list of all policies. The current state of each scheduled policy is indicated in the **Actions** column. If the **Disable** link is displayed, the policy's schedule is enabled. If the **Enable** link is displayed, the policy's schedule is disabled.
2. For the policy that you want to enable or disable, click **Enable** or **Disable**.



Note: If neither the **Enable** nor the **Disable** link appears, verify that an associated job is not currently running, and ensure that the SyncIQ module is licensed and enabled on the cluster.

Modify a SyncIQ policy

Before you modify an existing SyncIQ policy, consider how the changes will affect your cluster configuration and performance.



Note: You cannot modify a SyncIQ policy if its associated job is currently running.



Important: Modifying any of the following settings in an existing SyncIQ policy results in a *full* resynchronization of all affected data the next time the associated job runs:

- Any source-cluster directories, including the root path and any explicitly included or excluded directories

- Target path
- File criteria

Modifying any other SyncIQ policy settings, including most target-cluster settings, does not cause a full resynchronization to occur. However, if you specify a different physical target cluster or a new target path, a full resynchronization will occur, even if the new cluster's name is identical to the previous cluster's name.

1. On the **File System** menu, point to **SyncIQ**, and then click **Policies**.
The SyncIQ **Policies** page appears, and displays a list of all policies.
2. In the list of existing policies, click the name of the policy that you want to modify.
The SyncIQ **Edit Policy** page appears.
3. Modify the policy settings as needed, and then click **Submit**.

Delete a SyncIQ policy

Before permanently deleting a SyncIQ policy, consider how the deletion will affect your cluster configuration and performance. As an alternative to permanently deleting a policy, you can instead disable the policy, which keeps it available for future use.

If you delete a SyncIQ policy, the target association breaks and the target reverts to a read/write state. In addition, if you delete a SyncIQ policy, any reports for that policy are also deleted.

1. On the **File System** menu, point to **SyncIQ**, and then click **Policies**.
The SyncIQ **Policies** page appears, and displays a list of all configured policies.
2. In the **Actions** column for the policy that you want to delete, click **Delete**.
A dialog box appears, prompting you to confirm the deletion.
3. Click **Yes**.

SyncIQ policy and job operations

After you have created one or more SyncIQ policies, you can manually run SyncIQ jobs, view the status of any currently running jobs, or assess SyncIQ policies without running their associated jobs.

In addition to controlling SyncIQ operations from the source cluster, you can control certain SyncIQ operations from the target cluster. From the target cluster, you can cancel a currently running job, or you can break the association between a policy and its specified target directory by removing the associated cookie from the source. Breaking this association causes the source to fully resynchronize during the next job run.

A maximum of five SyncIQ jobs can run concurrently on a source cluster. You can start more than five jobs, or schedule more than five jobs to run during the same time period. However, only the first five jobs will run immediately; the remaining jobs wait in a queue until it is their turn to start. Note, however, that an unlimited number of jobs can run concurrently against a target cluster.

View SyncIQ job status

You can view the status of any currently running, or recently completed, SyncIQ jobs for which the current cluster is the source.

1. On the **File System** menu, point to **SyncIQ**, and then click **Summary**.
The SyncIQ **Summary** page appears, and displays a list of all currently running and recently completed SyncIQ jobs.
2. In the **Currently Running** area, review the information about any currently running SyncIQ jobs:
 - **Run:** Indicates the status (running or paused) of the job.
 - **Policy:** Displays the name of the associated SyncIQ policy. You can view or edit the policy's settings by clicking the policy name.
 - **Started:** Indicates the time at which the job started.
 - **Elapsed:** Indicates how much time has elapsed since the job started.

- **Transferred:** Indicates the number of files that have been transferred so far during the job run, and the total size of all transferred files.
- **Sync Type:** Indicates the type of synchronization being performed. The possible values are **Initial**, which indicates the full synchronization that is performed when a policy's job runs for the first time; **Upgrade**, which indicates a policy-conversion synchronization that occurs after upgrading the OneFS operating system or merging policies; and **Incremental**, which indicates a subsequent job that runs after a policy's full initial synchronization has been performed.
- **Source:** Displays the root path on the source.
- **Target:** Displays the target path.
- **Actions:** Displays any job-related actions that you can perform. You can perform an action by clicking its associated link. The specific links that appear depend on the job's current state.

3. In the **Recently Completed** area, review the information about any recently completed SyncIQ jobs:

- **Run:** Indicates the status of the job. A green icon indicates that the last job completed successfully. A yellow icon indicates that the last job did not complete successfully, but that an earlier job did complete successfully. A red icon indicates that jobs have run but that none of the jobs completed successfully. If no icon appears, the job has not run.
- **Policy:** Displays the name of the associated SyncIQ policy. You can view or edit the policy's settings by clicking the policy name.
- **Started:** Indicates the time at which the job started.
- **Ended:** Indicates the time at which the job finished running.
- **Duration:** Indicates the total duration of the job.
- **Transferred:** Indicates how many files that were transferred during the job run, and the total size of all transferred files.
- **Sync Type:** Indicates the type of synchronization that was performed. The possible values are **Initial**, which indicates the full synchronization that is performed when a policy's job runs for the first time; **Upgrade**, which indicates a policy-conversion synchronization that occurs after upgrading the OneFS operating system or merging policies; and **Incremental**, which indicates a subsequent job that runs after a policy's full initial synchronization has been performed.
- **Source:** Displays the root path on the source. This information may be unavailable for policies that were created in earlier versions of SyncIQ.
- **Target:** Displays the target path. This information may be unavailable for policies that were created in earlier versions of SyncIQ.
- **Actions:** Displays any job-related actions that you can perform. You can perform an action by clicking its associated link. You can click **View details** to view a more detailed report about the job run. The **SyncIQ Report Details** page appears, and displays detailed information about the job run and the parent policy, any reported error data, and other diagnostic details that can be useful when troubleshooting SyncIQ issues.

View SyncIQ local-target policy and job status

You can view information about, and the current status of, any SyncIQ policies and jobs for which the current, local cluster is the target.

1. On the **File System** menu, point to **SyncIQ**, and then click **Local Targets**.
The **SyncIQ Local Targets** page appears, and displays a list of any SyncIQ policies and currently running jobs for which the local cluster is the target.
2. Review the status information:



Note: Some data may be unavailable for policies that were created in earlier versions of SyncIQ.

- **Run:** Indicates the last known status of the job, as reported by the source cluster. A green icon indicates that the last job completed successfully. A yellow icon indicates that the last job did not complete successfully, but that an earlier job did complete successfully. A red icon indicates that jobs have run but that none of the jobs completed successfully. If no icon appears, the job has not run.



Note: A yellow or red icon may indicate that the policy has fallen into an unrunnable state. You can view more detailed policy-status information and, if necessary, resolve the source-target association, through the web administration interface on the source cluster.

- **Policy:** Displays the name of the associated SyncIQ policy. You can view or modify the policy's settings through the web administration interface on the source cluster.
- **Updated:** Indicates when data about the policy or job was most recently collected from the source cluster.
- **Source Cluster:** Displays the name of the source cluster.
- **Target Path:** Displays the path for the local target directory.
- **Coordinator IP:** Displays the IP address of the node on the source cluster that is acting as the job coordinator.
- **Actions:** Displays any job-related actions that you can perform. You can perform an action by clicking its associated link. The specific links that appear depend on the job's current state. The possible job actions are **Cancel** and **Break**.

Assess a SyncIQ policy

Before running a SyncIQ policy for the first time, you can assess the policy without actually running its associated job or transferring any files. This can be useful if, for example, you want to preview the size of the data set that would be affected if the policy's job were to run. During the policy-assessment process, the file system tree is surveyed at the source, and statistics are gathered for the files that would be subject to the data-synchronization action.

Prerequisite: Before you can assess a SyncIQ policy, you must enable the policy.



Note: You can assess a policy before it runs for the first time. However, once a job for that policy has actually been run, the assess-policy option is no longer available.

During the policy-assessment process, no bandwidth is consumed because the system does not connect to workers on the target and no data is transferred. Regardless of any policy options that may be configured, no data is modified or deleted on the target during the assessment process, nor is a snapshot created on the target cluster.

1. On the **File System** menu, point to **SyncIQ**, and then click **Policies**.
The SyncIQ **Policies** page appears, and displays a list of all configured SyncIQ policies.

2. In the **Actions** column for the policy that you want to test, click **Assess**.

If the **Assess** link does not appear, enable the policy and then try again.

The assessment process runs in the background, and can take several minutes to several hours to complete, depending on the number of files that need to be evaluated. When the assessment is complete, the results appear on the SyncIQ **Reports** page, where an assessed policy is indicated by a **Transferred** value of **Assessment**.

Manually run a SyncIQ job

As an alternative to, or in addition to, scheduling SyncIQ jobs to run automatically, you can manually start a job anytime.

Prerequisite: Before you can manually run a SyncIQ job, you must enable the associated policy.

1. On the **File System** menu, point to **SyncIQ**, and then click **Policies**.
The SyncIQ **Policies** page appears, and displays a list of all configured SyncIQ policies.

2. In the **Actions** column for the policy that you want to run, click **Start**.

If the **Start** link does not appear, enable the policy and then try again.

The SyncIQ job runs in the background. You can view the status of a running job on the SyncIQ **Summary** page.

Control SyncIQ operations

Through the web administration interface on the source cluster, you can control the operations for any SyncIQ jobs that are currently running, and you can perform other policy-specific actions. Stopping, pausing, or resuming a job does not affect its associated policy's enabled or disabled state.



Note: The specific links that appear depend on the current state of the policy or job.

1. On the **File System** menu, point to **SyncIQ**, and then click **Policies**.
The SyncIQ **Policies** page appears, and displays a list of all configured SyncIQ policies.
2. In the **Actions** column for the policy whose job you want to control, click the following links as needed:

- **Cancel:** Stops the policy's job, if a job is currently running.



Note: Canceling a currently running job does not affect any pending jobs that are scheduled in the SyncIQ policy.

- **Start:** Starts a job.
- **Pause:** Pauses the job, if a job is currently running. You can pause a job for no longer than seven days; after seven days, the job is automatically canceled.
- **Resume:** Resumes a paused job. If you resume a paused job, the job starts running from the point at which it was paused.
- **Reset:** Deletes the existing source snapshot and performs a full resynchronization. This link appears only if a policy is in a healthy state.



Caution: Depending on the amount of data being synchronized or copied, a full resynchronization can take a very long time to complete.

- **Resolve:** Places the policy back into a runnable state. This can be useful if you have corrected problems in a policy that SyncIQ previously identified as being unrunnable. This link appears only if a policy is in an unrunnable state.

Control SyncIQ local-target operations

In addition to controlling SyncIQ jobs for which the current cluster is the source, you can cancel any currently running jobs for which the current, local cluster is the target, or you can manually break the association between the local target and a SyncIQ policy. Canceling a job does not affect its associated policy's enabled or disabled state.



Note: From the target cluster, you cannot pause a SyncIQ job for which the local cluster is the target; you can, however, cancel a currently running job from the target cluster, and you can pause a SyncIQ job from the source cluster. Canceling a currently running local-target job does not affect any pending jobs that are scheduled in the SyncIQ policy.

1. On the **File System** menu, point to **SyncIQ**, and then click **Local Targets**.
The SyncIQ **Local Targets** page appears, and displays a list of any SyncIQ policies and currently running jobs for which the local cluster is the target.



Note: If the association between a policy and the local target has been broken, the policy does not appear in the list. You can view more detailed policy-status information and, if necessary, resolve a source-target association, through the web administration interface on the source cluster.

2. In the **Actions** column for the job that you want to control, click the following links as needed:

The specific links that appear depend on the job's current state. Some options may be unavailable for policies that were created in earlier versions of SyncIQ.

- **Cancel:** Stops the job by sending a cancel command to the source. If you cancel a local-target job, you cannot restart that job from the target cluster. The **Cancel** link appears only if the job is currently running.
- **Break:** Breaks the association between the policy and its specified target directory by removing the associated cookie from the source. Breaking this association causes the source to fully resynchronize the next time the job runs. The **Break** link appears only if the job is not currently running.

If you want to reestablish the association between the source and target, you can do so by resetting or resolving the policy through the web administration interface on the source cluster.

SyncIQ performance

You can customize SyncIQ performance settings by configuring network-performance and file-operation rules for specific days and periods of time. You can also monitor performance and system operations by viewing system logs.

Network-performance settings control network traffic usage, and file-operation settings control CPU and drive usage.

View SyncIQ performance statistics

The Isilon IQ system captures performance data for all SyncIQ-generated network usage, and CPU and drive usage for SyncIQ file operations. Any network or file-operation limits that you have configured are also represented in the performance data.

1. On the **File System** menu, point to **SyncIQ**, and then click **Performance**.
The SyncIQ **Performance** page appears.
2. In the **Network Performance** or **File Operations Performance** area, specify the time period for which you want to view performance statistics, and then click **Show History**.

For example, you might want to view five days' worth of network performance data ending on February 7, 2010, at 2:30 PM.

The SyncIQ network-usage statistics (measured in bytes per second) or file-operation statistics (measured in files per second) for the specified time period appear in the graph.

View SyncIQ performance rules

You can review the current settings for any configured SyncIQ performance rules. Network rules control network usage, and file-operation rules control CPU and drive usage.

1. On the **File System** menu, point to **SyncIQ**, and then click **Performance**.
The SyncIQ **Performance** page appears. The **Network Rules** and **File Operation Rules** tables at the bottom of the page list all configured network and file-operation performance rules.
2. Review the rule settings. You can sort the table data by clicking the column headings.
 - **State:** Indicates whether the rule is enabled or disabled.
 - **Days:** Indicates the day or days on which the rule is in effect.
 - **Start** and **End:** Indicates the time of day during which the rule is in effect.
 - **Limit:** Indicates the maximum network usage allowed (in bits per second), or CPU and drive usage allowed for file operations (in files per second) for SyncIQ processes.
 - **Description:** Displays any descriptive comments about the rule.
 - **Actions:** Displays a list of available rule-related actions. To modify the rule, click **Edit**. To delete the rule, click **Delete**.

Create a SyncIQ network-usage rule

You can configure network-usage rules that control how much SyncIQ-generated network traffic the cluster can process during a specified time period. These rules can be useful if, for example, you want to limit the amount of SyncIQ traffic during other resource-intensive operations.

You can configure multiple network-usage rules to run at different times. For example, you might allow only a small amount of SyncIQ network traffic during peak business hours but allow unlimited SyncIQ network usage during non-peak hours.

1. On the **File System** menu, point to **SyncIQ**, and then click **Performance**.
The SyncIQ **Performance** page appears.
2. In the **Network Rules** section, click **Add rule**.
The **Edit Limit** dialog box appears.
3. Specify the rule's **Status**:
 - **Enabled**: Click to enable the rule.
 - **Disabled**: Click to disable the rule.
4. In the **Description** box, type a descriptive comment about the rule.
This setting is optional, but a description can be helpful if you will be managing multiple rules.
5. In the **Limit** box, type an integer that represents, in bits per second, the maximum allowed network usage for SyncIQ traffic.



Note: Some units of file size specified for the network-usage limit are displayed differently in the web administration interface than how you originally enter them. Entering a number that is equal to a corresponding value in a larger unit of measure, such as Gb/s or Tb/s, results in the larger unit being displayed. For example, if you enter a value of **1000 Kb/s**, the web interface displays that value as "1 Mb/s." This change occurs because OneFS internally records network-usage limits in Kb/s and then converts them into Kb/s, Mb/s, Gb/s, or Tb/s for display.

6. Select the **Days** check boxes as needed to specify the day or days of the week on which the rule will apply.
7. In the **Start** and **End** lists, specify the times of day when the rule will start and end.



Note: A rule cannot span across more than one calendar day. For example, a rule that starts at 7:00 PM and ends at 11:00 PM is valid; however, a rule that starts at 7:00 PM and ends at 1:00 AM on the following day is invalid.

8. Click **Submit**.
The new rule appears in the **Network Rules** list on the SyncIQ **Performance** page.

Create a SyncIQ file-operation rule

You can configure file-operation rules that control how much SyncIQ-generated CPU or drive usage the cluster allows during specified time periods. These rules can be useful if, for example, you want to limit the amount of CPU or drive usage for SyncIQ file operations during other resource-intensive operations.

You can configure multiple file-operation rules to run at different times. For example, you might allocate only a small amount of CPU or drive usage to SyncIQ file operations during peak business hours, but allocate unlimited CPU or drive resources to SyncIQ file operations during non-peak hours.

1. On the **File System** menu, point to **SyncIQ**, and then click **Performance**.
The SyncIQ **Performance** page appears.
2. In the **File Operation Rules** section, click **Add rule**.
The **Edit Limit** dialog box appears.
3. Specify the rule's **Status**:

- **Enabled:** Click to enable the rule.
- **Disabled:** Click to disable the rule.

4. In the **Description** box, type a descriptive comment about the rule.

This setting is optional, but a description can be helpful if you will be managing multiple rules.

5. In the **Limit** box, type the maximum allowed CPU or drive usage, in files per second, for SyncIQ file operations.
6. Select the **Days** check boxes as needed to specify the day or days of the week on which the rule will apply.
7. In the **Start** and **End** lists, specify the times of day when the rule will start and end.



Note: A rule cannot span across more than one calendar day. For example, a rule that starts at 7:00 PM and ends at 11:00 PM is valid; however, a rule that starts at 7:00 PM and ends at 1:00 AM on the following day is invalid.

8. Click **Submit**.

The new rule appears in the **File Operation Rules** list on the **SyncIQ Performance** page.

Modify a SyncIQ performance rule

Before modifying a SyncIQ performance rule, consider how the changes will affect cluster performance.

1. On the **File System** menu, point to **SyncIQ**, and then click **Performance**.
The **SyncIQ Performance** page appears.
2. In the **Network Rules** or **File Operation Rules** section, in the **Actions** column, click the **Edit** link for the rule that you want to modify.
The **Edit Limit** dialog box appears.
3. Modify the rule's settings as needed, and then click **Submit**.

Delete a SyncIQ performance rule

Before deleting a SyncIQ performance rule, consider how the deletion will affect cluster performance.



Note: Alternatively, you can temporarily disable a performance rule instead of permanently deleting it.

1. On the **File System** menu, point to **SyncIQ**, and then click **Performance**.
The **SyncIQ Performance** page appears.
2. In the **Network Rules** or **File Operation Rules** section, in the **Actions** column, click the **Delete** link for the rule that you want to delete.
A confirmation dialog box appears.
3. Click **Yes**.
The rule is permanently deleted.

SyncIQ reports

SyncIQ generates reports that contain detailed information about job operations.

By default, the cluster retains a maximum of 2000 SyncIQ reports.

If multiple instances of a policy's job fail, SyncIQ consolidates those failure reports into a single report. You can view these consolidated reports, and you can view detailed subreports for each instance of the failure.



Note: If you delete a SyncIQ policy, any reports that were generated for that policy are also automatically deleted.



Note: You cannot customize SyncIQ report content.

View SyncIQ reports

You can view reports that contain detailed information about completed SyncIQ operations. You can also view the results of a policy assessment.

SyncIQ reports are available for viewing until they are automatically deleted. You can specify, either per-policy or through a global default setting, how long SyncIQ reports are stored on the cluster before they are automatically deleted.



Note: The report data displayed for an initial synchronization job (or for a job associated with a policy that was created in an earlier version of SyncIQ) is slightly different from report data for all subsequent, incremental jobs.

1. On the **File System** menu, point to **SyncIQ**, and then click **Reports**.
The SyncIQ **Reports** page appears, and displays a list of available SyncIQ reports.
2. Optionally, filter the displayed results by specifying the following report criteria:
 - **Page:** Type a page number to view a specific page, or click the previous-page and next-page buttons to scroll through the pages. Each page displays a maximum of 100 reports.
 - **Filters:** In these lists, you can click options to filter reports by specific policy names (or **All policies**), specific time periods (or **Anytime**), and by a job's run status (or **All status**).
 - **Refresh** button: Click to reload the SyncIQ **Reports** page and display only the reports that match the specified filtering criteria.
3. Review the report data:
 - **Status:** Indicates the status of the last run of the job. A green icon indicates that the last job completed successfully. A yellow icon indicates that the last job did not complete successfully, but that an earlier job did complete successfully. A red icon indicates that jobs have run but that none of the jobs completed successfully. If no icon appears, the job has not run.
 - **Policy:** Displays the name of the job's associated policy. You can view or edit a policy's settings by clicking the policy name.
 - **Started, Ended, and Duration:** Indicates when the job started and ended, and the duration of the job.
 - **Transferred:** Indicates the total number of files that were transferred during the job run, and the total size of all transferred files. An assessed policy is indicated by a **Transferred** value of **Assessment**.
 - **Sync Type:** Indicates the type of synchronization that was performed. The possible values are **Initial**, which indicates the full synchronization that is performed when a policy's job runs for the first time; **Upgrade**, which indicates a policy-conversion synchronization that occurs after upgrading the OneFS operating system or merging policies; and **Incremental**, which indicates a subsequent job that runs after a policy's full initial synchronization has been performed.
 - **Source:** Displays root path on the source.
 - **Target:** Displays the target path.
 - **Actions:** You can click **View details** to view a more detailed report about the job run. The SyncIQ **Report Details** page appears, and displays detailed information about the job run and the parent policy, any reported error data, and other diagnostic details that can be useful when troubleshooting SyncIQ issues. If multiple instances of a policy's job fail, SyncIQ consolidates those failure reports into a single report. You can view detailed subreports for each instance of a failure by clicking **View subreport** for the specified failure in the **Subreports** section of the **Report Details** page.

Configure global default SyncIQ report settings

You can configure a global default setting that specifies how long SyncIQ reports are stored on the cluster before they are automatically deleted. This global default setting is applied to any new SyncIQ policies that you create. However, any settings that you configure for specific policies take precedence over the global default setting.



Note: Modifying the global default SyncIQ report setting does not affect existing SyncIQ policies. This global default setting applies to only SyncIQ policies that you create after configuring these global settings.

1. On the **File System** menu, point to **SyncIQ**, and then click **Settings**.
The SyncIQ **Settings** page appears.
2. In the **Report Settings** area, in the **Retain reports for** box, type an integer that represents (in **days, weeks, months, or years**) how long reports will be stored on the cluster before they are automatically deleted.
The default setting is **365 days**.
3. Click **Submit**.

SyncIQ events

The SyncIQ module generates alerts when certain events occur during a SyncIQ job. These event alerts can be helpful when troubleshooting SyncIQ issues.

View and interpret SyncIQ events

You can view any event-specific alerts that have been generated in response to currently running SyncIQ jobs, as well as any event alerts for jobs that have already run. You can also view alerts for any events that you have canceled.

1. On the **File System** menu, point to **SyncIQ**, and then click **Events**.
The SyncIQ **Events** page appears, and displays a list of any current and historical SyncIQ events.
2. Review the data that is displayed:
 - **Sev:** Indicates the event's severity level.
 - **Instance ID:** Displays the unique system-generated event identifier.
 - **Start Time:** Indicates the date and time when the event started.
 - **End Time:** Indicates the date and time when the event ended.
 - **Message:** Describes the event.
 - **Scope:** Indicates the node on which the event occurred.
 - **Type:** Indicates the type of event. A "C" indicates a coalesced event, in which multiple identical events are consolidated into one event alert.
 - **Actions:** Displays any event-related actions that you can perform.

Manage SyncIQ event alerts

You can cancel a SyncIQ event alert, or view more detailed information about an event.

If you want to remove all alerts related to a specific event, you can cancel the alert.

1. On the **File System** menu, point to **SyncIQ**, and then click **Events**.
The SyncIQ **Events** page appears, and displays a list of any current and historical SyncIQ event alerts.



Note: The list does not include any alerts that have expired, or that have been canceled.

2. Manage event alerts as needed:
 - To view more information about an event, in the **Events** area, click the **View** link in the alert's **Actions** column.
The **View Event** page appears, and displays detailed information about the event.
 - To end an event alert, click the **Cancel** link in the alert's **Actions** column.

The Isilon SmartQuotas module

The Isilon SmartQuotas 2.0 module is an optional quota-management tool that monitors and enforces administrator-defined storage limits. Through its use of accounting and enforcement quota limits, reporting capabilities, and automated notifications, SmartQuotas can manage storage utilization, monitor disk storage, and issue alerts when disk storage limits are exceeded.

The Isilon SmartQuotas 2.0 module replaces all earlier versions of SmartQuotas. If you are using an earlier version of SmartQuotas, you must upgrade to 2.0 before you can use the SmartQuotas module.

The SmartQuotas module requires a separate license. For additional information about the SmartQuotas module, or to activate the module for your Isilon IQ clustered storage system, contact your Isilon Systems sales representative.

SmartQuotas overview

The Isilon SmartQuotas module can provision, monitor, and report disk storage usage at the user, group, and directory levels, and can send automated notifications when storage limits are exceeded or are being approached. SmartQuotas also supports flexible reporting options that can help you analyze data usage statistics for your Isilon cluster.

Types of quotas

SmartQuotas supports accounting quotas, which monitor—but do not limit—disk storage, and enforcement quotas, which monitor and limit disk storage.

Accounting quotas, which monitor disk storage utilization, are useful for auditing, planning, or billing purposes. For example, using SmartQuotas accounting quotas, you can:

- Track the amount of disk space used by various users or groups in order to bill each entity for only the disk space used.
- Review and analyze reports that help you identify storage usage patterns, which you can use to define storage policies for the organization and educate end users of the file system about using storage more efficiently.
- Intelligently plan for capacity expansions and future storage needs.

Enforcement quotas include all of the functionality of accounting quotas, and also limit disk storage and support notifications. Using enforcement quotas, you can logically partition an Isilon IQ cluster in order to control or restrict how much storage a given user, group, or directory can use. For example, you can set hard or soft capacity limits to ensure that adequate space is always available for key projects and critical applications, and to ensure that users of the cluster do not exceed their allotted storage capacity.

Enforcement quotas support three types of thresholds:

Threshold type	Description
Hard	A limit that cannot be exceeded. If an operation, such as a file write, causes a quota target to exceed a hard quota, the operation fails, an alert is logged to the cluster, and a notification is issued to any specified recipients.
Soft	A limit that can be exceeded until a grace period has expired. When a soft quota is exceeded, an alert is logged to the cluster and a notification is issued to any specified recipients; however, data writes are permitted during the

Threshold type	Description
	grace period. If the soft threshold is still exceeded when the grace period expires, data writes fail, and a hard-limit notification is issued to any specified recipients.
Advisory	An informational limit that can be exceeded. When an advisory quota threshold is exceeded, an alert is logged to the cluster and a notification is issued to any specified recipients. Reaching an advisory quota threshold does not prevent data writes.

Optionally, you can deliver real-time email notifications to users, group managers, or administrators about impending or current quota usage limit violations. Enforcement quotas support the following types of notifications and reminders:

- Threshold exceeded
- Over-quota reminder
- Grace period expired
- Write access denied

Types of quota entities

SmartQuotas supports a variety of entity types to which quotas can apply:

Entity type	Description
Directory	A specific directory and, optionally, all subdirectories.
Specific user	A specific user. Any specific-user quotas that you configure take precedence over a default-user quota.
Specific group	All members of a specific group. Any specific-group quotas that you configure take precedence over a default-group quota.
Default user	All users that are not associated with specific-user quotas. Associating a user quota with a default user quota creates a linked quota.
Default group	All groups that are not associated with specific-group quotas. Associating a group quota with a default group quota creates a linked quota.

Disk usage calculation methods for quotas

For each quota you configure, you can specify whether data-protection overhead is included in future disk-usage calculations. Configure these settings carefully, as the specified disk-usage calculation method can significantly affect the amount of disk space that is available for your end users to write to.

If you include data-protection overhead in a quota's usage calculations, any future disk-usage calculations for the quota will include the total amount of space required to store files and directories, as well as any space required to accommodate your specified data-protection settings, such as parity or mirroring. Tracking this type of disk usage for a quota effectively reduces the amount of space available to a user or group. For example, consider a user who is restricted by a 40 GB quota that includes data-protection overhead in its disk-usage calculations. If your cluster is configured with a 2x data-protection level and the user writes a 10 GB file to the cluster, that file actually consumes 20 GB of space: 10 GB for the file, and 10 GB for the data-protection overhead. In this example, the user has reached 50 percent of his 40 GB quota by writing a 10 GB file to the cluster.

If you do not include data-protection overhead in a quota's usage calculations, any future disk-usage calculations for the quota will include only the space required to store files and directories. Any space required for the cluster's specified

data-protection settings are not included. Consider the same example user, who is now restricted by a 40 GB quota that does not include data-protection overhead in its disk-usage calculations. If your cluster is configured with a 2x data-protection level and the user writes a 10 GB file to the cluster, that file consumes only 10 GB of space: 10 GB for the file, and no space for the data-protection overhead. In this example, the user has reached 25 percent of his 40 GB quota by writing a 10 GB file to the cluster. This method of disk-usage calculation is typically recommended for most quota configurations.

SmartQuotas upgrades

Isilon SmartQuotas 2.0 replaces all earlier versions of SmartQuotas. Earlier versions of SmartQuotas are not supported in this release of the Isilon OneFS file system. If you are running an earlier version of SmartQuotas, you must upgrade to SmartQuotas 2.0 before you can use the SmartQuotas module.

When you upgrade from an earlier version of SmartQuotas to SmartQuotas 2.0, you are prompted to convert your existing SmartQuotas files to the SmartQuotas 2.0 format. If you proceed with the conversion, your existing SmartQuotas quota settings are automatically converted to a format that is supported by SmartQuotas 2.0; in most cases, you do not need to manually reconfigure settings or convert files. However, you may need to reconfigure your SmartQuotas report settings.



Note: After converting to SmartQuotas 2.0, any pre-existing generated reports are no longer available through the web administration interface. To access reports that were generated in earlier versions of SmartQuotas, manually find and open them (as CSV files) from the directory in which they were originally saved. By default, this directory is `/ifs/data/smartquotas`, although a different location may have been specified during SmartQuotas setup.

Upgrade to SmartQuotas 2.0

Isilon SmartQuotas 2.0 replaces all earlier versions of SmartQuotas. Earlier versions of SmartQuotas are not supported in this release of the Isilon OneFS file system. If you are running an earlier version of SmartQuotas, you must upgrade to SmartQuotas 2.0 before you can use the SmartQuotas module.

If you are upgrading from an earlier version of SmartQuotas, the first time you access the SmartQuotas 2.0 module, the **SmartQuotas License Required** page may appear. If this page appears, you must enter a valid SmartQuotas license key before you can proceed.

1. After uploading the license key and enabling the SmartQuotas 2.0 module, the first time you access the SmartQuotas 2.0 module, the **Convert SmartQuotas Configuration File** page appears. Perform one of the following steps:
 - To convert your existing SmartQuotas files to the SmartQuotas 2.0 format, click **Convert my previous configuration**, and then click **Next**.
 - To bypass the conversion process and instead manually create new SmartQuotas policies, click **Skip the conversion**, and then click **Next**. The **SmartQuotas** page appears and no existing SmartQuotas files are converted.
2. If you clicked **Convert my previous configuration** in the previous step, on the next page of the wizard click **Convert configuration**.
The conversion process begins, and a status message appears while your files are being converted. When the conversion process is complete, an **Operation complete** message appears.
3. Click **Finish**.
The **SmartQuotas** page appears.



Note: The **Convert SmartQuotas Configuration File** page appears only once, the first time you access the SmartQuotas 2.0 module. However, you can perform the conversion process at any time through the command-line interface.

Quota management

You can modify the default storage quotas, and you can create quota limits and restrictions that apply to specific users, groups, or directories.

Create an accounting quota

An accounting quota monitors, but does not limit, disk usage.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Add Quota**.
The SmartQuotas **Add Quota** page appears.
2. In the **Apply quota to** list, click to specify the type of entity to which the quota will apply.
3. If you are creating a quota that applies to a specific user or specific group, type the name (in the form of a UID or GID) of the user or group in the **User or group** box.
4. In the **Path** box, type the fully qualified path of the directory, user, or group, or click **Browse** to locate it.
5. Click one of the **Snapshots** options.
 - To exclude snapshot usage from disk-usage calculations, click **Quota does not include usage for snapshots**. This is the default setting.
 - To include snapshot usage in disk-usage calculations, click **Quota includes usage for snapshots**. Any snapshots that were taken before you create a quota are not included in the quota's disk-usage calculations.
6. In the **Quota Policy** area, click **Accounting Quota**.
7. Click one of the **Usage calculation method** options.
 - To exclude data-protection overhead from disk-usage calculations, click **User data only**. This is the default setting, and is typically recommended for most quota configurations.
 - To include data-protection overhead in disk-usage calculations, click **User data and data-protection overhead**. Note that enabling this option may reduce the amount of disk space available to end users, potentially preventing users from writing to the cluster.
8. Click **Submit**.

Create an enforcement quota

An enforcement quota monitors and limits disk usage. You can create enforcement quotas that use any combination of hard limits, soft limits, and advisory limits.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Add Quota**.
The SmartQuotas **Add Quota** page appears.
2. In the **Apply quota to** list, specify the type of entity to which the quota will apply.
3. To create a quota that applies to a specific user or specific group, type the name (in the form of a UID or GID) of the user or group in the **User or group** box.
4. In the **Path** box, type the fully qualified path of the directory, user, or group, or click **Browse** to locate it.
5. Click one of the **Snapshots** options.
 - To exclude snapshot usage from disk-usage calculations, click **Quota does not include usage for snapshots**. This is the default setting.
 - To include snapshot usage in disk-usage calculations, click **Quota includes usage for snapshots**. Any snapshots that were taken before you create a quota are not included in the quota's disk-usage calculations.
6. In the **Quota type** area, click **Enforcement Quota**.
Additional options appear.
7. Configure threshold-specific enforcement limits as needed.

You must configure at least one threshold. Leaving any of the other threshold boxes blank creates an accounting-only quota for that threshold type, causing disk usage to be monitored but not limited.



Caution: Do not specify threshold settings that exceed the total size of the cluster.

- To configure a hard threshold that immediately prevents data writes when it is reached, type a **Hard threshold** integer value in **bytes, KB, MB, GB, TB, or PB**.
 - To configure a soft threshold that sends notifications and prevents data writes only after a grace period is reached, type a **Soft threshold** integer value in **bytes, KB, MB, GB, TB, or PB**. In the **Grace period** box, type an integer that specifies (in **minutes, hours, days, or weeks**) the duration of the grace period.
 - To configure an advisory threshold that sends notifications but never prevents data writes when it is reached, type an **Advisory threshold** integer value in **bytes, KB, MB, GB, TB, or PB**.
8. If you are configuring a directory quota with a hard threshold, specify a **Show disk capacity as** setting, which determines how much disk capacity is reported to end users.
 - To display only the size of the hard threshold to users, click **Size of hard threshold**.
 - To display the full size of the cluster to users, click **Size of cluster**.
 9. Click one of the **Usage calculation method** options.
 - To exclude data-protection overhead from disk-usage calculations, click **User data only**. This is the default setting, and is typically recommended for most quota configurations.
 - To include data-protection overhead in disk-usage calculations, click **User data and data-protection overhead**. Note that enabling this option may reduce the amount of disk space available to end users, potentially preventing users from writing to the cluster.
 10. Click **Submit**.
The **Quota** page appears.
 11. In the **Notifications** section, click the type of notification rules (**Global Default, None, Custom-basic, or Custom-advanced**) to use for the quota, and then click **Set**.



Note: To temporarily disable an enforcement quota instead of deleting it, change its quota type to **Accounting**. Toggling a quota's type between **Enforcement** and **Accounting** alternately enables and disables the quota while keeping its other settings intact.

Search for a quota

You can search for a quota using a variety of search criteria. This is a required step for many quota management tasks, such as modifying, cloning, or deleting quotas.



Note: Alternatively, you can view a list of all configured quotas on the SmartQuotas **Summary** page. On the **File System** menu, point to **SmartQuotas**, and then click **Summary**.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Edit Quotas**.
The SmartQuotas **Edit Quotas** page appears.
2. In the **Quota type** list, click to specify the entity type of the quota that you want to search for.
3. If you are searching for a **Specific user** or **Specific group** quota, type a full or partial user or group name in the **User or group** box.
You can include the ***** wildcard character in the **User or group** box.
4. Type a full or partial path in the **Path** box, or click **Browse** to locate it.
You can include the ***** wildcard character in the **Path** box.
5. If you want to expand the search to include the path's subdirectories, select the **Include subdirectories** check box.
6. If you are searching for a **Specific user** or **Specific group** quota, you can include linked quotas in the search results by selecting the **Include linked quotas** check box.

7. In the **View** list, click to specify the type of quota data to display in the search results:
 - To display threshold information in the quota search results, click **Thresholds**.
 - To display usage data in the search results, click **Usage**.
8. Click one of the **Include quotas...** options to specify the type of data displayed in the usage search results.
 - To display quotas both with and without data-protection overhead, click **Include quotas with and without overhead**.
 - To display only quotas with data-protection overhead, click **Include quotas with overhead**.
 - To display only quotas without data-protection overhead, click **Include quotas without overhead**.
9. Click **Search**.

Any quotas that match the specified search criteria appear in the search results area. The search returns a maximum of 100 quotas. If necessary, you can narrow your search criteria to display fewer results.
10. Review the search results:
 - An accounting quota, or an enforcement quota with a threshold value of zero, is indicated by dashes (--).
 - To sort the list of search results by a specific attribute, such as quota type or threshold type, click the corresponding column heading.

View quota settings

You can view a summary of all configured quotas, and you can view an individual quota's configured settings.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Summary**.

The SmartQuotas **Summary** page appears.
2. Review the SmartQuotas summary information in the **Quotas** area, which indicates the number of existing quotas of each entity type.



Note: The **Quotas** section also includes a **Import quotas** link that you can click in order to import quotas.

3. Search for the quota whose settings you want to view.



Note: Alternatively, on the SmartQuotas **Summary** page, you can click the **Edit** link to the right of an entity type. The SmartQuotas **Edit Quotas** page appears, and displays a list of all configured quotas of that entity type. To view an individual quota's settings, click that quota's **User or Group** link. The SmartQuotas **Quota** page appears, and displays the quota's settings in the **Settings** area. If you want to edit the quota's settings, click **Edit** in the **Settings** area. The SmartQuotas **Edit Quota** page appears.

4. In the search results list, click the name of the quota that you want to view.

The **Quota** page appears, displaying the current settings for the quota.
5. Review the quota's settings.

Not all of these fields appear in all types of quota summaries.

- **Type:** Indicates the type of entity to which the quota applies. If the quota is linked to a default quota, you can click a link in this field to display the default quota's settings.
- **User or group:** If the quota applies to a user or group, this field indicates the name of the user or group to which the quota applies.
- **Path:** Indicates the path to which the quota applies.
- **Snapshots:** Indicates whether snapshot disk usage is included in disk-usage calculations.
- **Hard threshold:** If the quota is used for enforcement, this field indicates the hard threshold value, if one is configured.
- **Soft threshold:** If the quota is used for enforcement, this field indicates the soft threshold value and grace-period setting, if a soft quota is configured.

- **Advisory threshold:** If the quota is used for enforcement, this field indicates the advisory threshold value, if one is configured.
- **Edit link:** Click to modify the quota's settings. This link does not appear if the quota is linked to a default quota.
- **Clone button:** Click to copy the quota.
- **Unlink from default button:** Click to unlink the quota from the default quota. This button appears only if the quota is linked to a default quota.
- **Delete button:** Click to delete the quota. This button does not appear if the quota is linked to a default quota.
- **Notifications:** Indicates the current notification setting for the quota.

Modify a quota

Before you modify a quota, consider how the changes will affect your file system and end users.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Edit Quotas**.
The SmartQuotas **Edit Quotas** page appears.
2. Search for the quota that you want to copy.
3. In the search results list, in the **User or Group** column, click the name of the quota that you want to modify.
The **Quota** page appears, displaying the current settings for the quota.
4. Click **Edit**.
The **Edit Quota** page appears.
5. Modify the quota's settings as needed.
6. Click **Save**.
The **Quota** page appears, displaying the quota's modified settings.
7. If you are configuring an enforcement quota, modify **Notifications** settings as needed, and then click **Set**.

Clone a quota

As an alternative to creating a new quota, you can create a new quota based on an existing quota's settings. The original quota's settings are copied to the new quota, enabling you to save time by modifying only the necessary settings.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Edit Quotas**.
The SmartQuotas **Edit Quotas** page appears.
2. Search for the quota that you want to copy.
3. In the search results list, click the name of the quota that you want to copy.
The **Quota** page appears, displaying the settings for the quota.
4. Click **Clone**.
The **Add Quota** page appears, with the applicable fields prepopulated with the original quota's settings.
5. Modify the new quota settings as needed.
6. Click **Submit**.
The **Quota** page appears, displaying the settings for the new quota.
7. If you are configuring an enforcement quota, modify **Notifications** settings as needed, and then click **Set**.

Delete a quota

Before you delete a quota, consider how the deletion will affect your file system and end users.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Edit Quotas**.
The SmartQuotas **Edit Quotas** page appears.
2. Search for the quota that you want to delete.
3. In the search results list, click the name of the quota you want to delete.
The **Quota** page appears, displaying the current settings for the quota.

4. Click **Delete**.
A dialog box appears, prompting you to confirm the deletion.
5. Click **OK**.

Unlink a quota from a default quota

If a quota is linked to a default quota, you cannot modify the linked quota's settings. You can, however, modify the default quota's settings or unlink the linked quota from the default quota. Unlinking a quota from a default quota creates a new, cloned quota that is based on the default quota's settings.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Edit Quotas**.
The SmartQuotas **Edit Quotas** page appears.
2. Search for the quota that you want to unlink from a default quota.
3. In the search results list, click the name of the quota that you want to unlink.
The **Quota** page appears, displaying the current settings for the quota and indicating the name of the default quota to which the quota is linked.
4. Click **Unlink from default**.
The **Edit Quota** page appears, displaying prepopulated settings that are based on the default quota settings.
5. Modify the now-unlinked quota settings as needed.
6. Click **Save**.

Import a SmartQuotas 2.0 configuration file

You can import quota configuration data in XML format. SmartQuotas 2.0 configuration files use a different XML format than earlier SmartQuotas versions. The steps for importing a SmartQuotas configuration file vary depending on whether the configuration file was created in SmartQuotas 2.0 or an earlier version of SmartQuotas, such as 1.0.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Summary**.
The SmartQuotas **Summary** page appears.
2. In the **Quotas** section, click **Import quotas**.
The **Import Quotas** page appears.
3. In the **Quota configuration file** box, type the fully qualified path of the file to import, or click **Browse** to locate it.
4. Click **Import**.

Import a SmartQuotas 1.x configuration file

You can import quota configuration data in XML format. SmartQuotas 2.0 configuration files use a different XML format than earlier SmartQuotas versions. The steps for importing a SmartQuotas configuration file vary depending on whether the configuration file was created in SmartQuotas 2.0 or an earlier version of SmartQuotas, such as 1.0.

When you upgrade to SmartQuotas 2.0 from an earlier version, you are prompted during the upgrade process to convert your existing SmartQuotas files to the SmartQuotas 2.0 format. However, you can perform the conversion at any time through the command-line interface.

1. Verify that the cluster is running SmartQuotas 2.0, and not an earlier version.
2. At the command prompt, run the following command:

```
isi quota import --convert-from-v1
```

The system locates any SmartQuotas 1.x configuration files and then converts them to the SmartQuotas 2.0 format.

Export a quota configuration file

You can export the results of a quota search query in XML format. This can be useful if, for example, you want to reuse your cluster's quota settings on another Isilon cluster or if you want to store quota configurations in another location outside of the cluster.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Edit Quotas**.
The SmartQuotas **Edit Quotas** page appears.
2. Perform a search query that matches the result set that you want to export.
3. When the search result set that you want to export appears on the **Search Quotas** page, click **Export**.
A dialog box appears, prompting you to save or open the `IsilonQuotaConfiguration.xml` file.
4. Save the XML file to your local file system.
The default file name is `IsilonQuotaConfiguration.xml`. You can modify the file name as needed.

Quota notifications

You can enable and configure default email notifications that are sent to users or groups when they are approaching, or have exceeded, a quota limit.

Enforcement quotas support the following notification settings. Note that any given quota can use only one of these settings.

Notification setting	Description
Global default	Uses the global default notification for the specified type of quota.
None	Disables all notifications for the quota.
Custom—basic	Enables you to create basic custom notifications that will apply to only this quota. You can configure basic notifications for any or all of the threshold types (hard, soft, or advisory) for the specified quota.
Custom—advanced	Enables you to create advanced custom notifications that will apply to only this quota. You can configure advanced notifications for any or all of the threshold types (hard, soft, or advisory) for the specified quota.

If you enable custom notifications that apply to a specific quota, you can optionally configure specific notifications for any or all of the threshold types (hard, soft, or advisory) for the specified quota, as described in the following table.

Condition	Action	Quota types
Quota exceeded	One-time instant notification (for example, "You have exceeded your quota.").	<ul style="list-style-type: none"> • Hard • Soft • Advisory
Over-quota reminder	One-time or recurring scheduled notifications (for example, "You are still over the quota. Please take action.").	<ul style="list-style-type: none"> • Hard • Soft • Advisory

Condition	Action	Quota types
Grace period expired	One-time or recurring scheduled notifications (for example, “Your grace period has expired. Please take action”).	<ul style="list-style-type: none"> Soft
Write access denied	One-time instant notification (for example, “You were denied writes because you are over the quota.”).	<ul style="list-style-type: none"> Hard Soft

Configure default global notification settings

You can configure default global quota notification settings that apply to all quotas of a specified threshold type. Any custom-notification settings you configure for a specific quota take precedence over these default global notification settings.

- On the **File System** menu, point to **SmartQuotas**, and then click **Notifications**.
The SmartQuotas **Notifications>Global Default** page appears, and displays the currently configured global default notification settings for hard thresholds.
- If you want to configure global default notification settings for a different threshold type, click the **Edit advisory threshold notifications** or **Edit soft threshold notifications** link.
The corresponding configuration fields appear.
- Specify a **Notification options** setting:
 - Use same settings for all events:** Click to use the same settings for all threshold-related events. This is the recommended setting.
 - Specify settings for each event:** Click to configure different settings for each type of threshold-related event. Additional options appear. Configure settings as needed, click **Submit**, and do not complete the rest of this procedure.
- To send threshold notifications to the owner of the entity, select the **Notify owner** check box.
- To send threshold notifications to a specific recipient, select the **Notify other** check box and then type the recipient’s **Email address**.
- To log an alert to the cluster, select the **Send alert** check box.
- Specify the notification’s **Interval** in days, weeks, months, or years.
- Specify the notification’s **Frequency**:
 - Once:** Click to send a notification only once, and then specify the time at which the notification will be sent.
 - Multiple times:** Click to issue recurring, scheduled notifications, and then specify how often notifications will be sent, and when the first and last notifications will be sent.
- Click **Submit**.

Configure basic custom quota notification settings

As an alternative to applying the default global quota notification settings to a quota, you can create custom notifications for specific quotas. Any custom-notification settings you configure for a specific quota take precedence over any default global notification settings.

- On the **File System** menu, point to **SmartQuotas**, and then click **Edit Quotas**.
The SmartQuotas **Edit Quotas** page appears.
- Search for the quota for which you want to configure a custom notification.
- In the search results list, click the name of the quota that you want to configure.
The **Quota** page appears, displaying the current settings for the quota.

4. Under **Notifications**, in the **Notification settings** list, click **Custom—basic**, and then click **Set**.
5. Click the **All thresholds** link.
The SmartQuotas Notifications **Per Quota** page appears.
6. To send threshold notifications to the owner of the entity, in the **Notifications** area, select the **Notify owner** check box.
7. To send threshold notifications to a specific recipient, select the **Notify other** check box and then type the recipient's **Email address**.
8. To log an alert to the cluster, select the **Send alert** check box.
9. Specify the notification's **Interval** in days, weeks, months, or years.
10. Specify the notification's **Frequency**:
 - **Once**: Click to send a notification only once, and then specify the time at which the notification will be sent.
 - **Multiple times**: Click to issue recurring, scheduled notifications, and then specify how often notifications will be sent, and when the first and last notifications will be sent.
11. Click **Submit**.

Configure advanced custom quota notification settings

The basic SmartQuotas notification settings are typically sufficient for most SmartQuotas deployments. However, if your configuration requires more granular or unique notification rules that apply to the each specific type of notification, you can configure a notification's advanced settings.

These settings apply to the default global notifications and to custom notifications.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Edit Quotas**.
The SmartQuotas **Edit Quotas** page appears.
2. Search for the quota for which you want to configure a custom notification.
3. In the search results list, click the name of the quota that you want to configure.
The **Quota** page appears, displaying the current settings for the quota.
4. Under **Notifications**, in the **Notification settings** list, click **Custom—advanced**, and then click **Set**.
5. Click to specify the type of threshold for which you want to configure a custom advanced notification (**Hard threshold**, **Soft threshold**, or **Advisory threshold**). The SmartQuotas Notifications **Per Quota** page appears.
6. In the **Notification options** area, click **Specify settings for each event**.
7. Specify advanced options as necessary.

Not all of these options appear on all versions of the page. The options that appear depend on the specified type of threshold.

Item	Description
Email message template settings	<p>Specify the type of email template to use for formatting email notifications. In the Email message template list, click one of the following options:</p> <ul style="list-style-type: none"> • Default: Uses the default email template. • Custom: Uses a custom email template that you specify. In the Template file box, type the full path of the email template (in .txt format), or click Browse to locate it. To preview a sample email notification with your custom template applied, click Preview.

Item	Description
Threshold exceeded settings	<p>Specify the types of notifications and alerts to issue when a hard, soft, or advisory threshold is exceeded.</p> <ul style="list-style-type: none"> • Notify owner check box: Select to send an email notification to the entity's owner. • Notify other check box: Select to send an email notification to another recipient, and then type the recipient's Email address. • Send alert check box: Select to log an alert message to the cluster. • In the Send a notification every box, type an integer that specifies a send-notification frequency in Hours, Days, or Weeks.
While over threshold settings	<p>Specify the types of notifications and alerts to issue when a hard, soft, or advisory threshold is exceeded.</p> <ul style="list-style-type: none"> • Notify owner check box: Select to send an email notification to the entity's owner. • Notify other check box: Select to send an email notification to another recipient, and then type the recipient's Email address. • Send alert check box: Select to log an alert message to the cluster. • Interval: Specify the notification and alert recurrence interval in days, weeks, months, or years. • Frequency: Specify the notification and alert frequency. <ul style="list-style-type: none"> • Once: Click to issue a notification or alert once, and then specify the time at which the notification or alert will be issued. • Multiple times: Click to run recurring, scheduled notifications or alerts, and then specify when they will begin and end.
Notification email template - grace period expiration date settings	<p>Specify the type of email template to use for formatting email notifications that are sent when a soft threshold's grace period has expired. In the Email message template list, click one of the following:</p> <ul style="list-style-type: none"> • Default: Use the default email template. • Custom: Use a custom email template that you specify. In the Template file box, type the full path of the email template (in .txt format), or click Browse to locate it. To preview a sample email notification with your custom template applied, click Preview.
Grace period expired reminder settings	<p>Specify the types of notifications and alerts to issue when a soft threshold's grace period has expired.</p> <ul style="list-style-type: none"> • Notify owner check box: Select to send an email notification to the entity's owner.

Item	Description
	<ul style="list-style-type: none"> • Notify other check box: Select to send an email notification to another recipient, and then type the recipient's Email address. • Send alert check box: Select to log an alert message to the cluster. • Interval: Specify the notification and alert recurrence interval in days, weeks, months, or years. • Frequency: Specify the notification and alert frequency. <ul style="list-style-type: none"> • Once: Click to issue a notification or alert once, and then specify the time at which the notification or alert will be issued. • Multiple times: Click to run recurring, scheduled notifications or alerts, and then specify when they will begin and end.
Write access denied settings	<p>Specify the types of notifications and alerts to issue when a hard threshold is exceeded, or when a soft threshold's grace period has expired.</p> <ul style="list-style-type: none"> • Notify owner check box: Select to send an email notification to the entity's owner. • Notify other check box: Select to send an email notification to another recipient, and then type the recipient's Email address. • Send alert check box: Select to log an alert message to the cluster. • In the Send a notification every box, type an integer that specifies a send-notification frequency in Hours, Days, or Weeks.

8. Click **Submit**.

Map an email notification for a quota

Email notification mappings control how email addresses are resolved when the cluster sends a quota notification. If necessary, you can remap the domain used for SmartQuotas email notifications. You can remap Active Directory Windows domains, local UNIX domains, or both.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Email Notification Mapping**.
The SmartQuotas **Email Notification Mapping** page appears.
2. In **Windows** area or the **UNIX** area, in the applicable **Map to** box, type the name of the domain that you want to map email notifications to.
You can repeat this step as needed, if you want to map more than one domain.
3. To test the mapping or mappings, type a valid **Username** and then click **Test**.
A test message is sent to the specified email recipient.
4. Click **Submit**.

Configure a custom email notification template

If email notifications are enabled, you can specify templates for formatting email notifications. You can use the default email templates that are included with SmartQuotas, or you can create and upload your own email templates, using a combination of human-readable text and supported SmartQuotas variables.

SmartQuotas includes two default email templates: one for limit-exceeded notifications, and another for grace-period-initiation notifications.

If the default email notification templates do not meet your needs, you can configure and upload your own custom email templates for SmartQuotas notifications. An email template can contain human-readable plain text and, optionally, variables that represent specific values. You can use any, all, or none of the available SmartQuotas variables in your templates. Template files must be saved in `.txt` format.

1. Create a file in `.txt` format that includes any combination of human-readable text and SmartQuotas variables. The supported variables are described in the following table.

Variable	Description	Human-readable example
<i>ISI_QUOTA_PATH</i>	Path of quota domain	/ifs/data
<i>ISI_QUOTA_THRESHOLD</i>	Threshold value	20 GB
<i>ISI_QUOTA_USAGE</i>	Disk space currently in use	10.5 GB
<i>ISI_QUOTA_OWNER</i>	Name of quota domain owner	jsmith
<i>ISI_QUOTA_TYPE</i>	Threshold type	Advisory
<i>ISI_QUOTA_GRACE</i>	Grace period, in days	5 days
<i>ISI_QUOTA_EXPIRATION</i>	Expiration date of grace period	Fri Feb 23 14:23:19 PST 2007

2. Save the template file to an appropriate directory on the Isilon cluster, and then, through the web administration interface, specify when and how to use the template file.

Example: Quota-exceeded custom template

The following example illustrates a custom email template that is used to notify recipients about an exceeded quota.

Text-file contents

```
The disk quota on directory <ISI_QUOTA_PATH> owned by
<ISI_QUOTA_OWNER> was
exceeded.
The <ISI_QUOTA_TYPE> quota limit is <ISI_QUOTA_THRESHOLD>, and
<ISI_QUOTA_USAGE> is currently in use. Please free some disk space
by deleting
any unnecessary files.
For more information, contact Jane Anderson in IT.
```

Human-readable email contents

```
The disk quota on directory /ifs/data/sales_tools/collateral owned
by djohnson
was exceeded.
The hard quota limit is 10 GB, and 11 GB is currently in use. Please
free some
disk space by deleting any unnecessary files.
For more information, contact Jane Anderson or Joe Brown in IT.
```

Quota reports

The SmartQuotas module supports flexible reporting options that enable you to more effectively manage cluster resources and analyze usage statistics. This section describes how to view reports and customize report settings that are used to track disk usage.



Note: If you have upgraded from SmartQuotas 1.x to SmartQuotas 2.0, any pre-existing generated reports are no longer available through the web administration interface. To access reports that were generated in SmartQuotas 1.x, manually find and open them (as CSV files) from the directory in which they were originally saved. By default, this directory is `/ifs/data/smartquotas`, although a different location may have been specified during SmartQuotas setup.

Configure quota report settings

You can configure quota report settings that track disk usage. These settings determine whether and when scheduled reports are generated, and where and how reports are stored.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Report Settings**.
The SmartQuotas **Report Settings** page appears.
2. Specify a **Run scheduled reports** setting:
 - To enable scheduled reports, click **Enable**. If you enable scheduled reports, the reports will automatically run according to the schedule you specify.
 - To disable scheduled reports, click **Disable**. If you disable scheduled reports, you can still run live, unscheduled reports at any time.
3. Specify the scheduled-report **Interval** in days, weeks, months, or years.
4. Specify the scheduled-report **Frequency**:
 - **Once**: Click to run a report once, and then specify the time at which the report will run.
 - **Multiple times**: Click to run recurring, scheduled reports, and then specify when the reports will begin and end.
5. In the **Keep** box, type an integer that represents the maximum number of scheduled reports to store.
This value determines the maximum number of scheduled reports that are available for viewing on the SmartQuotas **Reports** page. When the maximum number of reports has been stored, the system deletes the oldest reports to make space for new reports as they are generated.
6. In the **Report directory** box, type the full path of the directory in which you want to store scheduled reports, or click **Browse** to locate it.
7. Under **Live Reports**, in the **Keep** box, type an integer that represents the maximum number of live, on-demand reports to store.
This value determines the maximum number of live reports available for viewing on the SmartQuotas **Reports** page. When the maximum number of reports has been stored, the system deletes the oldest reports to make space for new reports as they are generated.
8. In the **Report directory** box, type the full path of the directory in which you want to store live reports, or click **Browse** to locate the directory.
9. Click **Submit**.

View quota reports

You can view quota reports in order to monitor and analyze disk storage utilization.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Reports**.

The SmartQuotas **Reports** page appears and displays a list of all reports available for viewing. The number of reports displayed is determined by the **Keep** values that you configured on the SmartQuotas **Report Settings** page.

2. In the **Report Date** column, click the date link for the report that you want to view.
The **View Quota Report** page appears, displaying the report for the specified date.
3. Specify report criteria as needed:
 - a. In the **Quota type** list, click the type of entity to generate a report for.
 - b. If viewing a user or group report, in the **User or group** box, type the UID or GID of the entity.
 - c. In the **Path** box, type the full path of the entity directory, or click **Browse** to locate it.
 - d. To include linked quotas in the report results, select the **Include linked quotas** check box.
 - e. To include subdirectories in the report results, select the **Include subdirectories** check box.
 - f. To display only storage-overage violations in the report results, select the **Show violations only** check box.
4. Click **Search**.
Any report data that matches the search criteria appears in a list.
5. Review the report results:
 - **Applies To:** Indicates the type of entity to which the report applies. If the report applies to a specific user or group, the appropriate UID or GID also appears in this field.
 - **Path:** Indicates the path to which the report applies.
 - **Usage:** Indicates the amount of disk storage currently being used in B, KB, MB, GB, TB or PB.
 - **Usage w/Overhead:** Indicates the total amount of disk storage currently being used, including all disk space in use for overhead operations.
 - **File count:** Indicates the number of files in use.
 - **Threshold:** Indicates the amount of disk storage allocated to the specified entity in B, KB, MB, GB, TB or PB.
 - **% Used:** Indicates the percentage of allowed disk storage being used.
 - **Overage:** If an overage has been reported, this column indicates the amount of the overage in B, KB, MB, GB, TB or PB.



Note: You can sort the report results by threshold type by clicking the **Hard thresholds**, **Soft thresholds**, or **Advisory thresholds** link.



Note: To download a copy of the report as a comma-separated value (CSV) file, click **Download**.

Run a live quota report

In addition to running regular, scheduled quota reports, you can run live, on-demand quota reports at any time in order to view quota usage statistics.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Quota Usage**.
The SmartQuotas **Quota Usage** page appears.
2. Specify your report criteria:
 - a. In the **Quota type** list, click the type of entity to generate a report for.
 - b. If generating a user or group report, in the **User or group** box, type the UID or GID of the entity.
 - c. In the **Path** box, type the full path of the entity directory, or click **Browse** to locate it.
 - d. To include linked quotas in the report results, select the **Include linked quotas** check box.
 - e. To include subdirectories in the report results, select the **Include subdirectories** check box.
 - f. To display only storage-overage violations in the report results, select the **Show violations only** check box.
3. Click **Search**.
Any reports that match the search criteria appear in a list on the **Live Report** page.
4. Review the report results:

- **Applies to:** Indicates the type of entity to which the report applies. If the report applies to a specific user or group, the appropriate UID or GID appears in this field.
 - **Path:** Indicates the path to which the report applies.
 - **Usage:** Indicates the amount of disk storage being used in B, KB, MB, GB, TB or PB.
 - **Usage w/Overhead:** Indicates the total amount of disk storage currently being used, including all disk space in use for overhead operations.
 - **File Count:** Indicates the number of files in use.
 - **Threshold:** Indicates the amount of disk storage allocated to the specified entity in B, KB, MB, GB, TB or PB.
 - **% Used:** Indicates the percentage of allowed disk storage being used.
 - **Overage:** If an overage has been reported, this column indicates the amount of the overage in B, KB, MB, GB, TB or PB.
5. If necessary, you can further refine the report results, or you can sort the report results according to threshold type by clicking the **Hard thresholds**, **Soft thresholds**, or **Advisory thresholds** link.
 6. If you want to save the quota report to the cluster, click **Save Report**.
The report appears in the list of saved reports on the SmartQuotas **Reports** page.
 7. If you want to download the report data in comma-separated value (CSV) format, click **Download**.

Download a quota report

You can download quota reports as comma-separated values (CSV) files, which can be useful for storing report information outside of the cluster or importing it into other applications.



Note: You cannot transfer quota reports between Isilon clusters.

1. On the **File System** menu, point to **SmartQuotas**, and then click **Reports**.
The SmartQuotas **Reports** page appears and displays a list of all reports available for viewing. The number of reports displayed is determined by the **Keep** values that you configured on the SmartQuotas **Report Settings** page.
2. Search to display the report results that you want to download.
3. When the report results that you want to download appear on the SmartQuotas **Reports** page, click **Download**.
A file-download dialog box appears, prompting you to open or save the `IsilonQuotaReport.csv` file.
4. Click **Save** to save the CSV file to your local file system, or click **Open** to open the file in another application, such as Microsoft Excel.

The Isilon SnapshotIQ module

The Isilon SnapshotIQ module is an optional software application that captures point-in-time images of data stored on an Isilon cluster.

SnapshotIQ is a flexible data protection application that is available as an option for the Isilon IQ cluster. SnapshotIQ provides a convenient local "insurance policy" for your users with a minimum of effort. You can schedule multiple snapshots at hourly, daily, weekly, monthly, or yearly recurring intervals, for directory, sub-directory, or file-system levels. You can also create snapshots manually.

You can use SnapshotIQ as a standalone tool to provide user-initiated file restoration and staging of exported content. You can use snapshots as the first line of defense for backing up and restoring data against accidental deletion and local data corruption. You can also use SnapshotIQ in conjunction with other OneFS features, such as backup and the SyncIQ replication module, to enhance the power and flexibility of those applications.

Snapshot schedule settings

With SnapshotIQ, you can establish multiple schedules for generating automatic snapshots at recurring intervals. SnapshotIQ's flexible scheduling enables you to configure snapshots with daily, weekly, monthly, and yearly frequencies. You can also specify naming patterns, durations, and expiration dates.



Note: After you create snapshot schedules, the system continues to generate snapshots (based on their recurrence patterns) until you remove them.

Snapshot schedule configuration

Before you can configure the frequency of specific snapshot schedules, you must first complete the initial, basic SnapshotIQ settings. After you complete the initial configuration, you can specify the frequency of the snapshot's schedule.

Configure basic SnapshotIQ schedule settings

Before you can create snapshot schedules, you must first configure the basic SnapshotIQ configuration steps. After you configure the basic SnapshotIQ settings, you can specify the frequency of the snapshot schedule.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Schedules**.
The SnapshotIQ **Schedules** page appears.

2. Click **Add schedule**.
The **Snapshot Schedule Basic Details** page appears.

3. In the **Schedule name** box, type a name for the snapshot schedule.
Snapshot names can be any valid file name that is not all numeric characters. Valid file names can contain a period or decimal point. However, a single or double period or decimal point (. or ..) alone is not a valid name. Forward-slash characters (/) are not permitted in file names.

4. In the **Snapshot pattern** box, type a snapshot naming pattern for automatically generated snapshots.
A snapshot pattern is a template that uses date variables to name generated snapshots with what is essentially a time-and-date stamp. It is recommended that you use the snapshot naming convention `filename_duration_variables`. Supported date variables include **%Y** for year, **%m** for month, **%d** for day, **%H** for hour, and **%M** for minute. You can separate variables using the dash, colon, and underscore characters. A pattern can be any string, but it is recommended that you configure the pattern to avoid non-unique snapshot names. For instance, a pattern of **sched_%d** would generate identical names after the first week of usage.

For example, if you want to create a weekly snapshot of data belonging to your organization's marketing department, you might create a snapshot pattern formatted like so: `marketing_weekly_%Y-%m-%d`. This would result in a snapshot named, for instance, `marketing_weekly_2008-08-31`. To generate more precise naming, you could add the hour and minute variables to the pattern.

5. In the **Path** box, type the full path of the directory that you want to capture in the snapshot, or click **Browse** to locate and select it.

The path can be any directory in the `/ifs` directory tree. The full path name is required, and wildcards are not allowed.

6. Optionally, in the **Alias** box, type an alias name for the snapshot schedule.

An alias is an alternative name that is usually shorter or easier to enter than the schedule name. Alias names have the same character restrictions as the schedule names. The alias points to the latest snapshot created by the schedule.

7. Click **Next**.

The **Snapshot Schedule: Scheduling Details** page appears.

After you complete these initial configuration steps, you can configure the following snapshots to occur on a daily, weekly, monthly, or yearly basis on the Snapshot Schedule **Scheduling Details** page.

Schedule a daily snapshot

You can configure your cluster to schedule a daily snapshot.

Prerequisite: Before you can specify a snapshot's frequency, you must first complete the initial schedule configuration of the snapshot.

1. On the **Snapshot Schedule: Scheduling Details** page, click **Daily**.
The page dynamically updates to display daily **Interval** options.
2. In the **Every** box, type the number of days that will elapse between snapshots.
3. In the **Every** list, select **days**(which includes both weekdays and weekends) or **weekdays**.
4. Under **Frequency**, select whether the snapshot will be generated **Once** or **Multiple times** on the days that match the daily criteria, and then specify the time when the snapshot will be generated in the hour and minute lists.
The time settings are based on a 24-hour clock, where 12:00 is noon and 00:00 is midnight.
5. Under **Expiration/Duration**, specify whether and when the snapshots produced by the schedule will be automatically deleted:
 - If you do not want snapshots created by this schedule to be automatically deleted, click **Never**.
 - If you want to automatically delete the snapshots after they reach a certain age, click **By elapsed age of snapshots**, type a positive integer in the **Retain snapshot for** box, and then specify a time interval in **hours**, **days**, **weeks**, **months**, or **years**.
6. Click **Next**.
The **Snapshot Schedule: Confirmation** page appears.
7. Review the snapshot schedule settings, and then click **Finish**.
The scheduled snapshot appears on the **Schedules** page.

Schedule a weekly snapshot

You can configure your cluster to schedule a weekly snapshot.

Prerequisite: Before you can specify a snapshot's frequency, you must first complete the initial schedule configuration steps for the snapshot.

1. On the **Snapshot Schedule: Scheduling Details** page, click **Weekly**.
The page dynamically updates to display weekly **Interval** options.
2. In the **Every** box, select the snapshot frequency interval in weeks, from 1 to 52 weeks.
3. Select the check box for the day of the week on which the snapshot will be generated. You also select multiple days.

4. Under **Frequency**, specify whether the snapshot will be generated **Once** or **Multiple times** on the days matching the weekly criteria, and then specify the time when the snapshot will be generated in the hour and minute lists.
The time settings are based on a 24-hour clock, where 12:00 is noon and 00:00 is midnight.
5. Under **Expiration/Duration**, specify whether and when the snapshots produced by the schedule will be automatically deleted:
 - If you do not want snapshots created by this schedule to be automatically deleted, click **Never**.
 - If you want to automatically delete the snapshots after they reach a certain age, click **By elapsed age of snapshots**, type a positive integer in the **Retain snapshot for** box, and then specify a time interval in **hours, days, weeks, months, or years**.
6. Click **Next**.
The **Snapshot Schedule: Confirmation** page appears.
7. Review the snapshot schedule settings, and then click **Finish**.
The scheduled snapshot appears on the **Schedules** page.

Schedule a monthly snapshot

You can configure your cluster to schedule a monthly snapshot.

Prerequisite: Before you can specify a snapshot's frequency, you must first complete the initial schedule configuration steps for the snapshot.

1. On the **Snapshot Schedule: Scheduling Details** page, click **Monthly**.
The page dynamically updates to display monthly **Interval** options.
2. Select one of the following monthly interval options:
 - Monthly by a specific day
 - Monthly by a relative day or weekday
 - Monthly by a specific date
3. Under **Frequency**, specify whether the snapshot will be generated **Once** or **Multiple times** on the days that match the monthly criteria, and then specify the time when the snapshot will be generated in the hour and minute lists.
The time settings are based on a 24-hour clock, where 12:00 is noon and 00:00 is midnight.
4. Under **Expiration/Duration**, specify whether and when the snapshots produced by the schedule will be automatically deleted:
 - If you do not want snapshots created by this schedule to be automatically deleted, click **Never**.
 - If you want to automatically delete the snapshots after they reach a certain age, click **By elapsed age of snapshots**, type a positive integer in the **Retain snapshot for** box, and then specify a time interval in **hours, days, weeks, months, or years**.
5. Click **Next**.
The **Snapshot Schedule: Confirmation** page appears.
6. Review the snapshot schedule settings, and then click **Finish**.
The scheduled snapshot appears on the **Schedules** page.

Schedule a yearly snapshot

You can configure your cluster to schedule a yearly snapshot.

Prerequisite: Before you can specify a yearly frequency, you must first complete the initial schedule configuration steps for the snapshot.

1. On the **Snapshot Schedule: Scheduling Details** page, click **Yearly**.
The page dynamically updates to display yearly **Interval** options.
2. Select one of the following yearly interval options:

- Every year on a specific month and day
 - Every year on a specific month and date
 - Every year on a relative day or weekday of a specific month
3. Under **Frequency**, specify whether the snapshot will be generated **Once** or **Multiple times** on the days that match the yearly criteria, and then specify the time when the snapshot will be generated in the hour and minute lists.
The time settings are based on a 24-hour clock, where 12:00 is noon and 00:00 is midnight.
 4. Under **Expiration/Duration**, specify whether and when the snapshots produced by the schedule will be automatically deleted:
 - If you do not want snapshots created by this schedule to be deleted, click **Never**.
 - If you want to delete the snapshots after they reach a certain age, click **By elapsed age of snapshots**, type a positive integer in the **Retain snapshot for** box, and then specify a time interval in **hours, days, weeks, months, or years**.
 5. Click **Next**.
The **Snapshot Schedule: Confirmation** page appears.
 6. Review the snapshot schedule settings, and then click **Finish**.
The scheduled snapshot appears on the **Schedules** page.

Snapshot schedule management

After you create a snapshot schedule, you can modify its settings, such as the frequency and interval of the schedule, and the name of and other information about the schedule. You can also delete a snapshot schedule if it is no longer needed.

Modify an existing snapshot schedule

After you create a snapshot schedule, you can view and modify its settings, such as the frequency and interval of the schedule, and the name of and other information about the schedule.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Schedules**.
The **Schedules** page appears.
2. Click the name of the schedule you want to modify.
The **Snapshot Schedule: Basic Details** page appears.
3. Modify the schedule's basic settings as needed, and then click **Next**.
The **Snapshot Schedule: Scheduling Details** page appears.
4. Review the scheduling details and modify the **Interval** and **Frequency** settings as needed.
5. Modify the **Expiration/Duration** settings as needed. If you click **By elapsed age of snapshots**, you must also review and modify the retention period.
6. Click **Next**.
The **Snapshot Schedule: Confirmation** page appears.
7. Review the snapshot settings, and then click **Finish**.
The modified snapshot appears on the **Schedules** page.

Delete a snapshot schedule

You can delete a snapshot schedule if it is no longer needed.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Schedules**.
The **SnapshotIQ Schedules** page appears.
2. To the right of the snapshot schedule that you want to delete, click **Delete**.
A confirmation page appears, prompting you to verify the deletion.
3. Click **Confirm**.
This removes the schedule from the **SnapshotIQ Schedules** page; it will no longer generate snapshots.

Manual snapshots

SnapshotIQ provides several methods for capturing cluster data and cluster state information. You can configure SnapshotIQ to automatically take snapshots according to a schedule that you specify. You can also take snapshots manually if you want to generate an immediate snapshot of the system, directory, or subdirectory.

Manually create a snapshot

You can configure your cluster to manually create a snapshot if you want to generate an immediate snapshot of the system, directory, or subdirectory.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Take Snapshot**.
The **Take Snapshot** page appears.
2. In the **Name** box, type a name for the manual snapshot.
Snapshot names can be any combination of characters that is not all numeric. Valid snapshot names can contain a period or decimal point. However, a single or double period or decimal point (. or ..) alone is not a valid name. Forward-slash characters (/) are not permitted in file names.
3. In the **Path** box, type the full path of the directory that you want to capture in the manual snapshot.
The path of the directory to be captured can be anywhere in */ifs* directory tree. The full path is required.
4. Optionally, in the **Alias** box, type an alias name.
An alias is an alternative name that is usually shorter or easier to enter than the schedule name. The alias points to the latest version of the manually created snapshot.
5. In the **Expiration** area, specify whether and when the manual snapshot will be automatically deleted:
 - To specify an expiration for the manual snapshot, select a date and time when the snapshot will automatically be deleted.
 - To keep the snapshot indefinitely, click **Never**.
6. Click **Submit**.
The **Existing Snapshots** page appears, with the newly created manual snapshot displayed.

Snapshot management

After snapshots have been created, you may need to find an individual snapshot and review its properties. In the SnapshotIQ application, you can view a list of pending scheduled snapshots, or the latest snapshots that have been generated, and select individual snapshots to review in detail.

View SnapshotIQ summary information

The **SnapshotIQ Summary** page provides an overview of the SnapshotIQ application. You can use it to review SnapshotIQ status, monitor recently completed and pending snapshots, and navigate to related SnapshotIQ pages to view, add, and modify schedules.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Summary**.
The **Summary** page appears.
2. On this page, you can perform the following tasks:
 - Under **Settings**, you can view the SnapshotIQ current status and settings, or click **Edit settings** to modify them.
 - Under **Recent Snapshots**, you can view the five most recently created snapshots, or click **View existing snapshots** to see all existing snapshots.
 - Under **Pending Snapshots**, you can view the next five pending snapshots, or click **View all pending snapshots** to see all scheduled snapshots.

- Using the links at the top of the page, you can perform all of the previous actions, as well as view snapshot schedules and take a manual snapshot.

View a recent snapshot

After snapshots have been created, you may need to find an individual snapshot and review its properties. You can view a list of snapshots that have been generated during a specific time period, and select individual snapshots to review in detail. The detailed snapshot view shows you properties such as creation date, path, size, and expiration date.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Existing Snapshots**.
The **Existing Snapshots** page appears.
2. In the **Show** list, specify whether you want to view snapshots that were taken in the past day, week, month, quarter, or year, or all existing snapshots. The default setting is **Snapshots for the past week**.
The page dynamically updates to show the names, creation dates, expiration dates, paths, file system information, and aliases of snapshots for the selected time period.
3. If you want to view or edit the settings for a specific snapshot, click its **Name** or **Alias**.
The **Edit Snapshot** page appears.

View a pending snapshot

In the SnapshotIQ application, multiple schedules may be in place to automatically make point-in-time captures of various cluster directories. You can view the dates and times of upcoming snapshots that will be generated by these schedules, as well as their recurrence and duration properties.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Pending Snapshots**.
The **Pending Snapshots** page appears.
2. In the **Show** list, specify whether you want to view snapshots that are scheduled to be taken in the next day, week, month, quarter, or year.
The page dynamically updates to show the schedule names, upcoming creation dates, and names of the snapshots for the selected time period.
3. If you want to view or edit the settings for a specific pending snapshot, click its name in the **Schedule** column.
The **Snapshot Schedule: Basic Details** page appears.

Rename or modify a snapshot

After snapshots are generated (either through a schedule or manually), you may want to change some properties. You can rename snapshots and modify their expiration dates.



Note: Usage issues relating to snapshot storage and deletion are complex. Deleting a snapshot (either manually or through expiration or duration settings) in order to free up space on the cluster might or might not result in recovering the same amount of space used by the snapshot. Deleting the oldest snapshot generally frees up the amount of space it was using, but deleting other snapshots might result in some or all of those disk blocks simply being transferred to another active snapshot.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Existing Snapshots**.
The SnapshotIQ **Existing Snapshots** page appears.
2. In the **Show** list, specify whether you want to view snapshots that were taken in the past day, week, month, quarter, or year, or all existing snapshots. The default setting is **Snapshots for the past week**.
The page dynamically updates to show the names, creation dates, expiration dates, paths, file system information, and aliases of snapshots for the selected time period.
3. In the **Name** column, click the snapshot you want to rename, modify, or delete.
The SnapshotIQ **Edit Snapshot** page appears.
4. Take one of the following actions:

- To give the snapshot a new name, click **Rename**, type a new name in the **Snapshot name** box, and then click **Submit**.
- To modify the snapshot's creation date, click **Edit expiration**, select a new expiration date and time, and then click **Submit**.
- To delete the snapshot, click **Delete**, and then click **Confirm**.

The SnapshotIQ **Edit Snapshot** page appears.

Delete a snapshot

After snapshots are generated (either through a schedule or manually), you may want to change some properties. You can rename snapshots, modify their expiration dates, or delete snapshots.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Existing Snapshots**.
The **Existing Snapshots** page appears.
2. In the **Show** list, specify whether you want to view snapshots that were taken in the past day, week, month, quarter, or year, or all existing snapshots. The default setting is **Snapshots for the past week**.
The page dynamically updates to show the names, creation dates, expiration dates, paths, file system information, and aliases of snapshots for the selected time period.
3. In the **Name** column, click the snapshot you want to delete.
The SnapshotIQ **Edit Snapshot** page appears.
4. Click **Delete**, and then click **Confirm**.
The SnapshotIQ **Edit Snapshot** page appears.

Snapshot usage and reserve settings

You can monitor the SnapshotIQ *usage*, which refers to the amount and percentage of file system space that is being used for individual snapshots, and the totals for all snapshots.

You can also view and configure the size of the SnapshotIQ *reserve*. Reserve is the percentage of file system space that is set aside exclusively for snapshot storage. Note that reserve is not a maximum boundary; in other words, snapshots can use more than the reserve. For example, if the SnapshotIQ reserve is set to 20 percent, normal user data is limited to 80 percent of the file system. Snapshots can use any amount, up to 100 percent of the available file system space, if necessary. The default level for reserve is 0 percent (that is, no space is exclusively set aside for snapshots).



Important: If you modify any SnapshotIQ settings (including changing the reserve level), Windows connections will be reset, even if you made no changes to the SMB configuration parameters. There might be a delay of up to two minutes for SMB to restart.

View individual snapshot usage

You can view the amount and percentage of file system space that is being used for individual snapshots.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Existing Snapshots**.
The **Existing Snapshots** page appears.
2. Review individual snapshot usage levels:
 - **Size** indicates the amount of file system space consumed by an individual snapshot.



Note: The percentage value does not include snapshot data shared with later snapshots or with the current active filesystem version. Usage issues relating to snapshot storage and deletion can be complicated. Deleting a snapshot (either manually or through expiration or duration settings) in order to free up space on the cluster might or might not result in recovering the same amount used by the snapshot. Deleting the

oldest snapshot generally frees up the amount of space it was using, but deleting other snapshots might result in some or all of those disk blocks simply being transferred to another active snapshot.

- **% Filesystem** shows the percentage of file system space used by an individual snapshot.
- **% Reserve** shows the percentage of space set aside exclusively for snapshot storage that has been consumed by an individual snapshot.

View overall snapshot usage

You can view the amount and percentage of file system space that is being used by all snapshots.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Summary**.
The **Summary** page appears.
2. Under **Settings**, review the following **Status** fields for SnapshotIQ usage information:
 - **Snapshot size** indicates the total amount of file system space used for snapshots.
 - **% of cluster capacity** shows the percentage of file system space used for snapshots.



Note: The percentage value does not include snapshot data that is shared with later snapshots or with the current active filesystem version. Usage issues relating to snapshot storage and deletion can be complicated. Deleting a snapshot (either manually or through expiration or duration settings) in order to free up space on the cluster might or might not result in recovering the same amount used by the snapshot. Deleting the oldest snapshot generally frees up the amount of space it was using, but deleting other snapshots might result in some or all of those disk blocks simply being transferred to another active snapshot.

- **Snapshot reserve** indicates the percentage of file system space set aside exclusively for snapshot storage.
- **Reserve used** shows the percentage of the snapshot reserve that has been consumed by snapshots.

Configure snapshot reserve levels

You can configure the size of the SnapshotIQ reserve, which is the percentage of file system space that is set aside exclusively for snapshot storage.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Summary**.
The **Summary** page appears, and displays the current percentage of space on the file system that is allocated exclusively for snapshot reserve and the amount that has been used by snapshots.



Note: Reserve is not a maximum boundary. In other words, snapshots can use more than the reserve. For example, if the SnapshotIQ reserve is set to 20 percent, normal user data is limited to 80 percent of the file system. Snapshots can use any amount, up to 100 percent of the available file system space, if necessary. The default level for reserve is 0 percent (that is, no space is exclusively set aside for snapshots).

2. Click **Edit settings**.
The SnapshotIQ **Settings** page appears.
3. Type a the percentage value in the **Reserve** box for the snapshot reserve amount, and then click **Submit**.



Note: If you set the SnapshotIQ reserve to a non-zero value (such as 1 percent, 5 percent, or 10 percent), it diminishes general storage availability, even if SnapshotIQ is currently disabled. Therefore, you must use caution when setting reserve values because this setting directly affects your total cluster storage capacity.

File and folder restoration

The SnapshotIQ application enables you to create snapshots of data and state information in a directory for safekeeping. If data is accidentally erased, lost, or otherwise corrupted or compromised, you can use the snapshot files to restore the data.

Windows users who have the Shadow Copy Client installed on their local computers can find snapshots through the **Previous Versions** tab for the folder's properties in Windows Explorer.



Note:

- To use the shadow copy emulation tool to view and restore files and folders that have been captured by SnapshotIQ, you may need to also download and install interface enhancements for Windows clients. Windows XP SP2, Windows 2003 Server, and Windows Vista versions include the shadow copy client feature by default; Windows XP versions earlier than SP2 do not.
- A shadow copy operation is limited to 64 previous versions, so exceeding that number can prevent access to all versions. This is true even though OneFS itself can support many more than 64 versions.

Restore a deleted file using Shadow Copy Emulation

You can use Shadow Copy Emulation to restore deleted files through Windows Explorer.

1. In Windows Explorer, navigate to the folder in which the deleted file was stored.
2. Position the cursor over a blank space in the folder, right-click, and then click **Properties**.
The **Properties** dialog box for the folder appears.



Note: If you rest the cursor over a file, that file will be selected, rather than the folder. For this procedure, you must select the folder, not an individual file.

3. Click the **Previous Versions** tab.
4. Select the folder version that contained the file before it was deleted and then click **View**.
5. Click the deleted file and then drag-and-drop or cut-and-paste the shadow copy file to the folder that it was deleted from, or to another location where you want it to be available.

Restore a corrupted or overwritten file using Shadow Copy Emulation

Restoring a corrupted or overwritten file is typically easier than recovering a deleted file, because you can right-click the file itself rather than the folder containing it.

1. In Windows Explorer, navigate to the folder that contains the corrupted or overwritten file.
2. Right-click on the file, and then click **Properties**.
The **Properties** dialog box appears.
3. Click the **Previous Versions** tab.
4. Select one of the existing previous versions of the file.




Note: To view the old version of the file, click **View**.


5. To copy the old version to another location, click **Copy**, and then select the new location.
6. To replace the current version of the file with the older version, click **Restore**.

Restore a deleted folder using Shadow Copy Emulation

You can use Shadow Copy Emulation to restore deleted folders through Windows Explorer.

 **Note:** A shadow copy operation is limited to 64 previous versions, so exceeding that number can prevent access to all versions. This is true even though OneFS itself can support many more than 64 versions.

1. In Windows Explorer, navigate to the folder in the deleted file was stored.
2. Position the cursor over a blank space in the folder, right-click, and then click **Properties**. The **Properties** dialog box appears.


 **Note:** If you rest the cursor over a file, that file will be selected, rather than the folder. For this procedure, you must select the folder, not an individual file.

3. Click the **Previous Versions** tab, and then perform the appropriate action:
 - To view a previous version, select the folder, and then click **View**. The folder contents appear.
 - To recover the full contents of that folder as well as all subfolders, click **Restore**.
 - To copy the old version of the folder to another location, click **Copy**, and then select the location.

SnapshotIQ settings


After you license and activate SnapshotIQ, you must configure the application before you can use it.

When you activate the SnapshotIQ module, the application is automatically enabled. You can disable or enable SnapshotIQ at any time.


 **Important:** If you modify any SnapshotIQ settings (including changing the reserve level), Windows connections will be reset, even if you made no changes to the SMB configuration parameters. There may be a delay of up to two minutes before SMB restarts.

Enable or disable SnapshotIQ

SnapshotIQ is enabled by default when the module is licensed. Disabling SnapshotIQ turns off the automatic creation or deletion of snapshots, and prevents the manual creation of snapshots.

 **Note:** SnapshotIQ is automatically enabled if it has been activated through the Licensing process.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Settings**. The **Settings** page appears.
2. Next to the **Service** option, click either **Enable** or **Disable** to change the status of SnapshotIQ, and then click **Submit**. The updated status for SnapshotIQ appears on the **Summary** page.

 **Note:** If you change the SnapshotIQ status from **Enable** to **Disable**, any scheduled snapshots will not be generated. Snapshots that are scheduled to be automatically deleted, based on expiration dates or durations, will not be deleted. However, all existing snapshots will be retained, and any snapshot schedules will still be shown in the **Schedules** list. You can still manually modify or delete snapshots.

Configure basic SnapshotIQ settings

After you activate the SnapshotIQ license, you must enable and configure the application before you can use it. You can configure basic functions for the SnapshotIQ application, including enabling automatic creation and automatic

deletion of snapshots, and setting the amount of space devoted exclusively to snapshot storage. You can also configure advanced settings that control user access and directory visibility.



Important: Whenever you change any SnapshotIQ settings (including changing the reserve level), Windows connections will be reset, even if you make no changes to the SMB configuration parameters. There might be a delay of up to two minutes for SMB to restart.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Settings**.
The **Settings** page appears.
2. For **Auto create state**, click **On** to enable the scheduled generation of snapshots, or click **Off** to disable scheduled snapshot creation.



Note: The ability to create manual snapshots is not affected by the **Auto create state** setting.

3. For **Auto delete state**, click **On** to enable the automatic deletion of generated snapshots, or click **Off** to disable automatic snapshot deletion.
4. In the **Reserve** box, type a percentage value for the disk space on the cluster that will be reserved for storing snapshots. The **Reserve** value is not a maximum boundary. In other words, snapshots can use more space than the reserve. For example, if the SnapshotIQ reserve is set to 20%, normal user data is limited to 80% of the file system; snapshots can use any amount, up to 100% of the available file system space, if necessary.



Important: The default value for the **Reserve** setting is 0%, meaning there is no space exclusively set aside for snapshots. If the SnapshotIQ **Reserve** value is set to any non-zero value, it diminishes general storage availability, even if the snapshots feature is currently disabled. Use caution when setting reserve values because they affect your total cluster storage capacity.

5. Click **Submit**.
The **Summary** page appears.

Configure advanced SnapshotIQ settings

You can configure advanced options for root directory and subdirectory access and visibility for NFS, Windows, and local users.



Important: Whenever you make any SnapshotIQ setting changes (including changing the reserve level), Windows connections will be reset, even if you made no changes to the SMB configuration parameters. There might be a delay of up to two minutes for SMB to restart.

1. On the **File System** menu, point to **SnapshotIQ**, and then click **Settings**.
The **Settings** page appears.
2. Click **View advanced settings**.
The page expands to display additional configuration settings.
3. Under **Global**, configure **Visibility & accessibility** to control the overall availability of the SnapshotIQ application to users on the cluster. If you want to make snapshots available to users, click **On**.



Note: Clicking **Off** prevents local, NFS, and SMB users from seeing the `.snapshot` directories. Disabling visibility and accessibility also deactivates all the option buttons below it for NFS, Windows, and local user accessibility and visibility.

4. Enable or disable the options under **NFS Settings**, **Windows Settings**, and **Local Settings** by clicking **On** or **Off** as appropriate for your cluster users:
 - **Root directory accessible:** To make the root directory accessible to users, click **On**. For local users the root directory is `/ifs/.snapshot`; for NFS users the root directory is the client mount point; and for Windows (SMB) users the root directory is the top level of the share they are connected to. To prevent user access, click **Off**. Clicking **Off** also toggles the **Root directory visible** and **Subdirectories accessible** options to **Off**.

- **Root directory visible:** To make snapshots in the `/ifs/.snapshot` directory visible to users in directory listings, click **On**. To hide them from users, click **Off**.



Note: Making the root directory visible may interfere with tree removal, and could affect folder deletion in Windows.

- **Subdirectories accessible:** To make the subdirectories in the `/ifs` directory accessible to users, click **On**. To prevent user access, click **Off**.

5. Click **Submit**.

The **Summary** page appears.

The Isilon iSCSI module

The Isilon iSCSI module is an optional application that provides Internet Small Computer System Interface (iSCSI) support on an Isilon cluster. The iSCSI module enables you to configure, provision, and connect block storage for Microsoft Windows, Linux, and VMware systems.

The iSCSI module requires a separate license. For additional information about the iSCSI module, or to activate the module for your Isilon clustered storage system, contact your Isilon Systems sales representative.

iSCSI overview

The Isilon iSCSI module enables you to create and manage iSCSI targets on the Isilon cluster, and enables initiators on client computers to send SCSI commands to those targets. Targets are exposed as SCSI block devices that can be used for storing unstructured or structured data.

The iSCSI module uses a client/server model, in which initiators on client computers make requests to targets that are located on storage nodes on an Isilon cluster. iSCSI targets contain one or more logical units, each uniquely identified by a logical unit number (LUN), which the client can format on the local file system and use as it would a physical disk device.

You can configure separate data protection levels for each logical unit, through Isilon FlexProtect or data mirroring.

For basic access control, you can configure each target to limit access to a specific list of initiators. You can also require initiators to authenticate themselves to a target by requiring Challenge-Handshake Authentication Protocol (CHAP) authentication.

The iSCSI module supports the following features:

- Microsoft Internet Storage Name Service (iSNS) server
- Microsoft Volume Shadow Copy Service (VSS)
- Isilon SmartConnect Advanced dynamic IP allocation
- Isilon FlexProtect
- 8x data mirroring
- LUN cloning
- One-way CHAP authentication
- Initiator access control

Security

The iSCSI module supports Challenge-Handshake Authentication Protocol (CHAP) and initiator access control for connections to individual targets.

These security options can be used together or separately:

- **CHAP authentication:** You can enable and configure CHAP authentication. If you enable CHAP authentication, iSCSI initiators must authenticate themselves by providing a valid CHAP user:secret pair that is configured in the target's CHAP secrets list.
- **Initiator access control:** By default, targets are open to all initiators. You can enable access control to mask a target to all initiators or to limit access to a specified list of allowed initiators. If a target is closed to a given initiator, that target will not be discoverable to the initiator, and any attempt to connect to the target will be denied to that initiator.

Supported iSCSI initiators

OneFS 6.5.0 and higher operating systems are compatible with the following iSCSI Initiators.

Operating System	iSCSI Initiator
Microsoft Windows 2003 (32-bit and 64-bit)	Microsoft iSCSI Initiator 2.08 and higher (Certified)
Microsoft Windows 2008 (32-bit and 64-bit)	Microsoft iSCSI Initiator (Certified)
Microsoft Windows 2008 R2 (64-bit only)	Microsoft iSCSI Initiator (Certified)
Red Hat Enterprise Linux 5	Linux Open-iSCSI Initiator (Supported)
VMware ESX 4.0 and ESX 4.1	iSCSI Initiator (Certified)
VMware ESXi 4.0 and ESXi 4.1	iSCSI Initiator (Certified)

Limitations and considerations

When planning your iSCSI deployment, be aware of the following limitations and considerations.

- Multipath I/O (MPIO) is recommended only for iSCSI workflows with primarily read-only operations, unless all MPIO sessions are connected to the same node. Performance decreases in proportion to the number of write operations because for each write request that comes into a node, the node must invalidate the data cache on all other nodes for the file that is being written to.
- The Isilon iSCSI module supports one-way Challenge-Handshake Authentication Protocol (CHAP) authentication; however, mutual CHAP authentication is not supported.
- The Isilon iSCSI module supports the importing of normal LUNs only; importing snapshot LUNs and shadow LUNs is not supported. You cannot back up and then restore a snapshot or shadow LUN, or replicate snapshot or shadow LUNs to another cluster. It is recommended that you deploy a backup application on the iSCSI client to back up iSCSI LUNs, in order to ensure that the LUN is in a consistent state at the time of backup.
- The Isilon iSCSI module does not support:
 - Internet Protocol Security (IPsec)
 - Multiple connections per session (MCS)
 - iSCSI host bus adaptors (HBAs)

Global iSCSI settings

You can enable or disable the iSCSI service and configure the iSNS client service, which is used by iSCSI initiators for target discovery. These settings are applied globally to all nodes in the cluster. You cannot modify these settings for individual nodes.

Monitor iSCSI sessions

If the iSCSI service is enabled on the cluster, each node listens on all configured network interfaces for iSCSI sessions. You can view a summary of current iSCSI sessions and throughput data on the cluster.



Note: If you want to view iSCSI throughput data, you must do so through the Isilon InsightIQ virtual appliance, which requires a separate license. For more information, contact your Isilon representative.

1. On the **File System** menu, point to **iSCSI**, and then click **Summary**.
The iSCSI **Summary** page appears.
2. Review the current throughput data and current session information.

- The **Current Throughput** section displays a chart that illustrates overall inbound and outbound throughput across all iSCSI sessions during the past hour, measured in kilobits per second (Kbps). This chart automatically updates every 15 seconds.
- The **Current Sessions** section displays information about each current connection between an initiator and a target, including the client and target IP addresses; node, target, and LUN; operations per second; and the inbound, outbound, and total throughput in bits per second. You can view details about a particular target by clicking the target name.

Configure the iSCSI service

You can enable or disable the iSCSI service for all nodes in the cluster.

Before you disable the iSCSI service, be aware of the following considerations:

- All current sessions will be terminated for all nodes in the cluster.
 - Initiators will not be able to establish new sessions until the iSCSI service is re-enabled.
1. On the **File System** menu, point to **iSCSI**, and then click **Settings**.
The iSCSI **Settings** page appears.
 2. In the **iSCSI Service** section, configure the service state as needed:
 - If the service is disabled, you can enable it by clicking **Enable**.
 - If the service is enabled, you can disable it by clicking **Disable**.

Configure the iSNS client service

You can configure, enable, or disable the Internet Storage Name Service (iSNS), which is used for target discovery by iSCSI initiators.

1. On the **File System** menu, point to **iSCSI**, and then click **Settings**.
The iSCSI **Settings** page appears.
2. In the **iSNS Client Service** section, configure the iSNS client service settings:
 - **iSNS server address**: Type the IP address of the iSNS server on which you want to register iSCSI target information.
 - **iSNS server port**: Type the iSNS server port number. The default port is **3205**.
3. Optionally, click **Test connection** to validate the iSNS configuration settings.
One of the following messages displays:
 - `Connection to iSNS server succeeded. The iSNS client service is configured.`
 - `Connection to iSNS server failed. Review the iSNS server address and iSNS server port and make any necessary changes.`
4. Click **Submit**.
You must complete the next step to enable the iSNS client service.
5. Optionally, modify the service state as needed:
 - If the service is disabled, you can enable it by clicking **Enable**. This allows communication between the cluster and the iSNS server.
 - If the service is enabled, you can disable it by clicking **Disable**. This prevents communication between the cluster and the iSNS server.

Configure default SmartConnect pools

You can configure a list of SmartConnect pools that iSCSI targets will connect through by default. Default SmartConnect pools are ignored by targets that already have specific SmartConnect pools configured.

Modifying the default SmartConnect pools does not affect existing iSCSI sessions.

1. On the **File System** menu, point to **iSCSI**, and then click **Settings**.
The iSCSI **Settings** page appears.
2. In the **Default SmartConnect Pools** section, click **Edit list**.
The **Manage SmartConnect Pools for Target** dialog box appears.
3. Optionally move pools between the **Available Pools** and **Selected Pools** lists by clicking a pool and then clicking the right or left arrow. To remove all selected pools at once, click **clear**.



Note: Valid SmartConnect pools require a subnet with a fully qualified domain name that points to a virtual IP address. Only valid SmartConnect pools will be configured.

4. Click **OK**.

iSCSI target management

iSCSI targets define connection endpoints and serve as container objects for logical units on an iSCSI server. iSCSI initiators on clients establish connections to those targets.

You can configure one or more targets for your iSCSI server, and each target can contain one or more logical units. You can control access to the target by configuring SmartConnect pools, initiator access control, and Challenge-Handshake Authentication Protocol (CHAP) authentication. The iSCSI module supports Microsoft Internet Storage Name Service (iSNS) server for target discovery.

Create a target

You can configure one or more iSCSI targets, each with its own initiator access control and authentication settings. Targets are required as container objects for logical units.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The iSCSI **Targets & Logical Units** page appears.
2. In the **Targets** section, click **Add target**.
The iSCSI **Add Target** page appears.
3. In the **Name** box, type a name for the target.
Valid names must begin with a letter and can contain only lowercase letters, numbers, and hyphens (-).
4. Optionally, in the **Description** box, type a descriptive comment for the target.
5. In the **Path** box, type the full path of the directory, beginning with **/ifs**, where logical unit number (LUN) directories will be created by default, or click **Browse** to locate a directory.



Note: This directory is used only if none is specified during LUN creation. The directory must be within the **/ifs** directory tree. The full path to the directory is required, and wildcard characters are not supported.

6. Optionally add one or more SmartConnect pools for the target to connect through. This setting will override any global default SmartConnect pools that are configured for iSCSI targets.
 1. For the **SmartConnect pool(s)** setting, click **Edit list**. The **Manage SmartConnect Pools for Target** dialog box appears.
 2. Optionally move pools between the **Available Pools** and **Selected Pools** lists by clicking a pool and then clicking the right or left arrow. To remove all selected pools at once, click **clear**.



Note: Valid SmartConnect pools require a subnet with a fully qualified domain name that points to a virtual IP address. Only valid SmartConnect pools will be added to the list.

3. Click **OK**.

7. Click **Submit**.

The iSCSI **Edit Target** page appears.

8. Optionally, under **Initiator Access Control**, enable and configure initiator access control settings.

a. Click **Enable** to restrict target access to initiators that are added to the initiator access list.

b. Click **Add initiator**.

The **Add Allowed Initiator** dialog box appears.

c. In the **Initiator name** box, type the name of the initiator that you want to allow access to, or click **Browse** to select from a list of provided or known initiators.

Valid initiator names must begin with the **iqn.** prefix.

d. Click **OK**.



Note: To continue adding initiators, click **OK and add another**. When you are finished adding initiators, click **OK**.

9. Optionally, under **CHAP Authentication**, enable and configure Challenge-Handshake Authentication Protocol (CHAP) settings.



Note: If CHAP authentication is enabled and the CHAP secrets list is empty, no initiators will be allowed to access the target.

a. Click **Enable** to require initiators to authenticate themselves to the target.

b. Click **Add username**.

The **Configure CHAP credentials** dialog box appears.

c. In the **Username** box, type the name that the initiator will use to authenticate against the target.

You can specify an initiator's iSCSI Qualified Name (IQN) as the username. Depending on whether you specify an IQN, valid usernames differ in the following ways:

- If you specify an IQN as the username, the **Username** value must begin with **iqn.** Valid characters allowed after the **iqn.** prefix can include alphanumeric characters, periods (.), hyphens (-), and colons (:).
- All other usernames can use alphanumeric characters, periods (.), hyphens (-), and underscores (_).



Note: CHAP usernames are case-sensitive.

d. In the **Secret** and **Confirm secret** boxes, type the secret that the initiator will use to authenticate against the target.

Valid CHAP secrets must be 12-16 characters long, and can contain any combination of letters, numbers, and symbols.

e. Click **OK**.

10. Click **Submit**.

The iSCSI **Targets & Logical Units** page appears, and the new target appears in the **Targets** section.

View target settings

You can view detailed information about a target, including its iSCSI Qualified Name (IQN), default logical unit number (LUN) directory path, total size (capacity), and SmartConnect pool settings, as well as associated logical units, initiator access control settings, and Challenge-Handshake Authentication Protocol (CHAP) authentication settings.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.

The iSCSI **Targets & Logical Units** page appears.

2. In the **Targets** section, click the target name.

The iSCSI **View Target** page opens and displays the following sections:

- **Target Details:** Displays the target name, IQN, description, default path, capacity, and SmartConnect pool settings. The name and IQN cannot be modified. To modify other settings, click **Edit target**.
- **Logical Units:** Displays any logical units that are contained in the target. You can add or import a logical unit, or manage existing logical units.
- **Initiator Access Control:** Displays the target's initiator access control status, and lists the names of any initiators that are allowed to access the target when access control is enabled. To enable or disable initiator access control, or to modify the list of initiators, click **Edit target** in the **Target Details** section.
- **CHAP Authentication:** Displays the target's CHAP authentication status, and lists the users for which user:secret pairs have been enabled. To enable or disable CHAP authentication, or to modify the list of CHAP secrets, click **Edit target** in the **Target Details** section.

Modify a target

You can modify a target's description, change the path where logical unit directories will be created, and modify the list of SmartConnect pools that the target will use. You can also manage the target's security settings, including initiator access control and authentication policy settings.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The iSCSI **Targets & Logical Units** page appears.
2. In the **Targets** section, under **Actions**, click **Edit** for the target that you want to modify.
The iSCSI **Edit Target** page appears.
3. Modify the target's settings as needed.



Note:

- Changing the default path does not affect existing logical units.
- Changing the security settings does not affect existing connections.

4. Click **Submit**.

Delete a target

When you delete a target, all logical unit numbers (LUNs) that are contained in the target are also deleted, and all LUN data is destroyed. In addition, any iSCSI sessions that are connected to the target are terminated. This operation cannot be undone.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The iSCSI **Targets & Logical Units** page appears.
2. In the **Targets** section, under **Actions**, click **Delete** for the target that you want to delete.
A confirmation dialog box appears.
3. Click **Yes**.
The iSCSI **Targets & Logical Units** page appears. The target, and all LUNs and LUN data that are contained in the target, are deleted, and any iSCSI sessions on the target are terminated.

Initiator access control

You can control which initiators are allowed to connect to a target by enabling initiator access control and configuring the target's initiator access list.

By default, initiator access control is disabled and all initiators are allowed to access the target. You can restrict access to a target by enabling access control and then adding initiators to the target's initiator access list. If you enable access control but leave the initiator access list empty, no initiators will be allowed to access the target.

Configure access control settings

You can configure access control settings to specify which initiators are allowed to connect to a target.

If initiator access control is enabled for an iSCSI target, access to that target is limited to a specified list of allowed initiators. Access control is disabled by default.



Note: Modifications to a target's access control settings are applied to subsequent connection requests. Current connections are not affected.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
2. In the **Targets** section, under **Actions**, click **Edit** for the target whose initiator access state you want to change.
The **iSCSI Edit Target** page appears.
3. In the **Initiator Access Control** section, configure the access control state.

- If access control is **disabled**, you can click **Enable** to restrict target access to initiators that are added to the initiator access list.



Note: If you disable access control and the initiator access list is empty, no initiators will be allowed to connect to the target. You can add initiators by clicking **Add initiator**.

- If access control is **enabled**, you can click **Disable** to grant all initiators access to the target.



Note: If you disable access control, the list of allowed initiators is ignored.

Add an initiator to the access list

You can control access to a target by adding initiators to its initiator access list. If you enable initiator access control, the initiator access list specifies which initiators are allowed to access the target.



Note: The initiator access list is ignored unless you enable initiator access control.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
2. In the **Targets** section, under **Actions**, click **Edit** for the target that you want to allow initiator access to.
The **iSCSI Edit Target** page appears.
3. In the **Initiator Access Control** section, click **Add initiator**.
The **Add Allowed Initiator** dialog box appears.
4. In the **Initiator name** box, type the name of the initiator that you want to allow access to, or click **Browse** to select from a list of provided or known initiators.
Valid initiator names require the **iqn.** prefix.
5. Click **OK**.

6. To continue adding initiators, click **OK and add another**.
7. When you are finished adding initiators, click **OK**.

Modify initiator settings

You can rename or replace an initiator that is allowed to connect to a target when access control is enabled.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
2. In the **Targets** section, under **Actions**, click **Edit** for the target whose initiator access list you want to modify.
The **iSCSI Edit Target** page appears.
3. In the **Initiator Access Control** section, under **Actions**, click **Edit** for the initiator that you want to modify.
The **Edit Allowed Initiator** dialog box appears.
4. Modify the initiator's settings as needed.
5. Click **OK**.

Remove an initiator from the access list

You can remove an initiator from a target's initiator access list so that the initiator will no longer be allowed to connect to a target if access control is enabled.



Note:

- If you remove all of the allowed initiators for a target and access control is enabled, the target will deny any new connections until you disable access control.
- Removing an allowed initiator for a particular target does not affect the initiator's access to other targets.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
2. In the **Targets** section, under **Actions**, click **Edit** for the target whose initiator access list you want to modify.
The **iSCSI Edit Target** page appears.
3. In the **Initiator Access Control** section, under **Actions**, click **Delete** for the initiator that you want to remove.
The **Confirm delete** dialog box opens.
4. Click **Yes**.
The **Confirm delete** dialog box closes, and the initiator is removed from the initiator access list.

CHAP authentication

The iSCSI module supports Challenge-Handshake Authentication Protocol (CHAP) for authenticating initiator connections to iSCSI targets.

You can restrict initiator access to a target by adding CHAP user:secret pairs to the target's CHAP secrets list and then enabling CHAP authentication. This requires initiators to provide a valid user:secret pair to authenticate their connections to the target. CHAP authentication is disabled by default.



Note: The Isilon iSCSI module does not support mutual CHAP authentication.

Enable or disable CHAP authentication

You can enable or disable CHAP authentication for individual targets.



Note: Modifications to a target's CHAP authentication status are applied to subsequent connection requests. Current connections are not affected.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
2. In the **Targets** section, under **Actions**, click **Edit** for the target whose CHAP authentication state you want to modify.
The **iSCSI Edit Target** page appears.
3. In the **CHAP Authentication** section, configure the initiator's CHAP authentication state.

- If CHAP authentication is **disabled**, you can click **Enable** to require initiators to authenticate themselves to the target.



Note: If CHAP authentication is enabled and the CHAP secrets list is empty, no initiators will be allowed to access the target. You can add CHAP user:secret pairs by clicking **Add username**.

- If CHAP authentication is **enabled**, you can click **Disable** to stop authenticating initiators to the target.




Note: If CHAP authentication is disabled, the CHAP secrets list is ignored.

Create a CHAP secret

To use CHAP authentication, you must create CHAP user:secret pairs in the target's CHAP secrets list. Initiators will be required to provide a valid user:secret pair to authenticate their connections to the target.



Note: CHAP secrets are used only if you enable CHAP authentication.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
 2. In the **Targets** section, under **Actions**, click **Edit** for the target that you want to create a CHAP secret for.
The **iSCSI Edit Target** page appears.
 3. In the **CHAP Authentication** section, click **Add username**.
The **Configure CHAP credentials** dialog box appears.
 4. In the **Username** box, type the name that the initiator will use to authenticate against the target.
You can specify an initiator's iSCSI Qualified Name (IQN) as the username. Depending on whether you specify an IQN, valid usernames differ in the following ways:
 - If you specify an IQN as the username, the **Username** value must begin with **iqn.** Valid characters allowed after the **iqn.** prefix can include alphanumeric characters, periods (.), hyphens (-), and colons (:).
 - All other usernames can use alphanumeric characters, periods (.), hyphens (-), and underscores (_).
-  **Note:** CHAP usernames are case-sensitive.
5. In the **Secret** and **Confirm secret** boxes, type the secret that the initiator will use to authenticate against the target.
Valid CHAP secrets must be 12-16 characters long, and can contain any combination of letters, numbers, and symbols.
 6. Click **OK**.

Modify a CHAP secret

You can modify the settings for a CHAP user:secret pair.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
2. In the **Targets** section, under **Actions**, click **Edit** for the target whose CHAP secrets list you want to modify.
The **iSCSI Edit Target** page appears.
3. In the **CHAP Authentication** section, under **Actions**, click **Edit** for the username of the CHAP secret that you want to modify.

The **Configure CHAP credentials** dialog box appears.

4. Modify settings as needed, and then click **OK**.

Delete a CHAP secret

You can delete a CHAP secret that is no longer needed.



Note: If you delete all of a target's CHAP secrets and CHAP authentication is enabled, no initiators will be allowed to access the target until you disable CHAP authentication.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
2. In the **Targets** section, under **Actions**, click **Edit** for the target whose CHAP secret you want to remove.
The **iSCSI Edit Target** page appears.
3. In the **CHAP Authentication** section, under **Actions**, click **Delete** for the CHAP secret that you want to remove.
The **Confirm delete** dialog box appears.
4. Click **Yes**.
The **Confirm delete** dialog box closes, and the CHAP user:secret pair is removed from the list of CHAP secrets.

iSCSI LUN management

A logical unit defines a storage object (a disk or disk array) that is made accessible by an iSCSI target on the Isilon cluster. Each logical unit is uniquely identified by a logical unit number (LUN). Although a LUN is an identifier for a logical unit, the terms are often used interchangeably.



Note: A logical unit must be associated with a target, and each target can contain one or many logical units.

The following table describes the three types of LUNs that the Isilon iSCSI module supports:

LUN type	Description
Normal	This is the default LUN type for clone LUNs and imported LUNs, and the only type available for newly created LUNs. Normal LUNs can be either writeable or read-only.
Snapshot	A snapshot LUN is a copy of a normal LUN or another snapshot LUN. Snapshot LUNs require little time and disk space to create, but they are read-only. You can create snapshot LUNs by cloning existing normal or snapshot LUNs, but you cannot create snapshot clones of shadow LUNs.
Shadow	A shadow LUN is a copy of a normal, snapshot, or shadow LUN. A shadow LUN is a compromise between a normal LUN and a snapshot LUN, and is implemented using overlay and mask files in conjunction with a snapshot. Shadow LUNs require very little time and disk space to create, and the LUN is fully writeable. You can create shadow LUNs by cloning or importing existing LUNs or by using the Microsoft Volume Shadow Copy Service (VSS) provider.

Create a logical unit

Each logical unit is uniquely identified by a logical unit number (LUN), which is assigned when you create the logical unit. You can configure a logical unit's target assignment, LUN value, LUN directory path, LUN size, provisioning policy, access state, write access, protection settings, and I/O optimization settings.

Follow these steps to create and configure a logical unit.



Note: When you create a logical unit, you must assign it to an existing iSCSI target. Each target can contain one or more logical units.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
2. In the **Logical Units** section, click **Add logical unit**.
The **iSCSI Add Logical Unit** page appears.
3. Optionally, in the **Description** box, type a descriptive comment for the logical unit.
4. In the **Target** list, click to specify a target to contain the logical unit.
5. Select one of the **LUN number** options.
 - To assign the next available number to the logical unit, click **Automatic**. This is the default setting.
 - To manually assign a number to the logical unit, click **Manual** and then, in the **Number** box, type an integer value. Valid values must be within the range 0-255, and must not be currently assigned to another logical unit within the target.

By default, the configured LUN value is used as part of the directory name that is created for storing the LUN data.

6. To manually specify the path where the LUN directory will be created, in the **Path** box, type the full path, beginning with **/ifs**, of the directory, or click **Browse** to locate the directory.
The directory must be within the **/ifs** directory tree. You must specify the full path to the directory, and wildcard characters are not allowed. The default path is **/ifs/iscsi/ISCSI.LUN.<TargetName>.<LUNnumber>**, where **<TargetName>** is the configured **Target** value and **<LUNnumber>** is the configured **LUN number** value.
7. In the **Size** box, specify the LUN capacity by typing an integer value and then selecting a unit of measure (**MB**, **GB**, or **TB**).
The minimum LUN size is 1 MB. The maximum LUN size is 32 TB. After you create a LUN, you can increase its size but you cannot decrease it.
8. Select one of the **Provisioning** options.
 - To specify that blocks are unallocated until they are written, click **Thin provision**. This is the default setting.
 - To immediately allocate all blocks, click **Pre-allocate space**.



Note: Allocation of all blocks for a large LUN can take hours or even days. Until the allocation is complete, the LUN will not be accessible to iSCSI initiators.

9. Select one of the **LUN access** options.
 - To make the LUN accessible, click **Online**. This is the default setting.
 - To make the LUN inaccessible, click **Offline**.
10. Select one of the **Write access** options.
 - To allow iSCSI initiators to write to the LUN, click **Read-Write**. This is the default setting.
 - To prevent iSCSI initiators from writing to the LUN, click **Read-Only**.
11. In the **Disk pool** list, click to specify a disk pool to contain the logical unit.
12. In the **SSD strategy** list, click to specify a strategy to use if solid-state drives (SSDs) are available.
 - **Metadata acceleration:** Creates a mirror backup of file metadata on an SSD and writes the rest of the metadata plus all user data on hard disk drives (HDDs). Depending on the global namespace acceleration setting, the SSD mirror may be an extra mirror in addition to the number required to satisfy the protection level.
 - **Avoid SSDs:** Never uses SSDs; writes all associated file data and metadata to HDDs only.
 - **Data on SSDs:** Similar to metadata acceleration, but also writes one copy of the file's user data (if mirrored) or all of the data (if not mirrored) on SSDs. Regardless of whether global namespace acceleration is enabled, any SSD blocks reside on the file's target pool if there is room. This SSD strategy does not result in the creation of additional mirrors beyond the normal protection level.
13. In the **Protection level** list, click to specify a protection policy for the logical unit. **Use iSCSI default** or one of the mirrored options (**2x - 8x**) is the recommended setting for LUNs.

14. Select one of the **SmartCache** options.

- To prevent write caching for files that contain LUN data, click **Disabled**. This is the recommended setting for LUNs.
- To allow write caching for files that store LUN data, click **Enable**.



Note: Write caching can improve performance, but can lead to data loss if a node loses power or crashes while uncommitted data is in the write cache. This can cause inconsistencies in any file system that is laid out on the LUN, rendering the file system unusable.

15. Select one of the **Data access pattern** options.

- To select a random access pattern, click **Random**. This is the recommended setting for LUNs.
- To select a concurrent access pattern, click **Concurrency**.
- To select a streaming access pattern, click **Streaming**. Streaming access patterns can improve performance in some workflows.

16. Click **Submit**.

The iSCSI **Targets & Logical Units** page appears.

View logical unit settings

You can view information about a logical unit, including its logical unit number (LUN), iSCSI target, LUN type, LUN directory path, iSCSI qualified name (IQN), and other settings.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.

The iSCSI **Targets & Logical Units** page appears.

2. In the **Logical Units** section, under **Target:LUN**, click the name of the logical unit that you want to view.

The iSCSI **View LUN** page appears and displays the following settings:

- **LUN:** Displays the numerical identifier of the logical unit. You can modify the LUN value by using the move operation.
- **Target:** Displays the name of the iSCSI target that contains the logical unit. You can modify the target by using the move operation.
- **Description:** Displays an optional description for the logical unit. You can modify the description by clicking **Edit LUN**.
- **Type:** Displays the LUN type (normal, shadow, or snapshot). This value cannot be modified.
- **Size:** Displays the LUN capacity. You can increase the size of normal or snapshot LUNs by clicking **Edit LUN**, but you cannot decrease the size. You cannot modify the size of snapshot LUNs.
- **Status:** Displays the connection status (online or offline) and write access permissions (read-only or read-write) of the LUN. You can modify write-access settings for normal or shadow LUNs by clicking **Edit LUN**. You cannot modify write-access settings for snapshot LUNs.
- **Path:** Displays the path to the directory where the LUN files are stored. You can change the path for normal or snapshot LUNs by using the move operation. You cannot modify the path for snapshot LUNs.
- **Disk pool:** Displays the disk pool of the LUN. You can modify the disk pool by clicking **Edit LUN**.
- **Protection level:** Displays the mirroring level (2x-8x) or FlexProtect protection policy for the LUN. You can modify the protection policy for normal or shadow LUNs by clicking **Edit LUN**. You cannot modify these settings for snapshot LUNs.
- **SmartCache:** Displays whether SmartCache is enabled or disabled. You can change this setting for normal or shadow LUNs by clicking **Edit LUN**. You cannot modify these settings for snapshot LUNs.
- **Data access pattern:** Displays the access pattern setting (Random, Concurrency, or Streaming) for the LUN. You can change the access pattern for normal or shadow LUNs by clicking **Edit LUN**. You cannot modify these settings for snapshot LUNs.
- **SCSI name:** Displays the iSCSI qualified name (IQN) of the LUN. You cannot modify this setting.

- **EUI:** Displays the extended unique identifier (EUI), which uniquely identifies the LUN. You cannot modify this setting.
- **NAA:** Displays the LUN's T11 Network Address Authority (NAA) namespace. You cannot modify this setting.
- **Serial number:** Displays the serial number of the LUN. You cannot modify this setting.

Control access to a logical unit

You can specify whether a logical unit is offline or online. In order for initiators to discover and connect to a logical unit, the logical unit must be online.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
2. In the **Logical Units** section, under **Actions**, click either **Offline** or **Online** for the logical unit whose access setting you want to modify.
The **Status** icon changes color to indicate whether the logical unit is offline (gray) or online (green).

Modify a logical unit

You can modify certain settings for a logical unit. Note that some logical-unit settings cannot be modified.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
2. In the **Logical Units** section, under **Actions**, click **Edit** for the logical unit that you want to modify.
The **iSCSI Edit Logical Unit** page appears.
3. Modify the logical unit's settings as needed.
4. Click **Submit**.

Delete a logical unit

Deleting a logical unit permanently deletes all data on the logical unit. This operation cannot be undone.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
2. In the **Logical Units** section, under **Actions**, click **Delete** for the logical unit that you want to remove.
A confirmation dialog box appears.
3. Click **Yes**.

Clone a logical unit

As an alternative to creating a new logical unit from scratch, you can create a logical unit that is based on an exact point-in-time copy of an existing logical unit. The original (source) logical unit's settings are copied to the new (clone) logical unit, enabling you to save time by modifying only the necessary settings.



Note: The cloning process requires snapshots; to use the cloning feature, the Isilon SnapshotIQ module must be enabled on the cluster. SnapshotIQ is a separately licensed feature. For information about purchasing a SnapshotIQ license, contact your Isilon sales representative.

Cloned copies of logical units can be part of the same target or a different target. Lun cloning, like LUN creation, is an asynchronous process. The status of in-progress cloning operations is monitored in the same way as LUN creation. A clone LUN is not accessible by iSCSI initiators until the cloning process is complete.

Depending on the clone LUN type, the contents (blocks) of a source LUN are either copied or referenced, and the attributes may or may not be copied. In general, snapshot and shadow clone operations are fast, whereas normal clones can take several minutes or even hours to create depending on the size of the LUN. The following table provides further details about each cloning operation:

Source LUN type	Clone LUN type	Result
Normal	Normal	A snapshot of the source LUN is created. The clone LUN is then created by copying the LUN data from the snapshot. After completing the copy, the snapshot is deleted. The copy process may take several hours to complete for large LUNs if the source LUN has a pre-allocated provisioning policy. The copy process may also take several minutes for thinly provisioned LUNs that are significantly utilized.
Normal	Snapshot	A snapshot of the source LUN is created. The clone LUN is configured to reference the data from the snapshot. The snapshot is deleted when the clone is deleted.
Normal	Shadow	A snapshot of the source LUN is created. The system then creates a shadow LUN that references data from the snapshot.
Snapshot	Normal	The clone LUN is created by copying the LUN data from the snapshot. The copy process may take several minutes to complete for large LUNs if the source LUN has a pre-allocated provisioning policy. The copy process may also take several minutes for thinly provisioned LUNs that are significantly utilized.
Snapshot	Snapshot	The clone LUN is configured to reference the data from the same snapshot that the source LUN references. The underlying snapshot will not be deleted when a LUN is deleted unless the LUN being deleted is the last LUN referencing the snapshot.
Snapshot	Shadow	The clone LUN is configured to reference the data from the same snapshot that the source LUN references. The underlying snapshot will not be deleted when a LUN is deleted unless the LUN being deleted is the only LUN referencing the snapshot.
Shadow	Normal	A snapshot of the source LUN is created. The clone LUN is then created by copying the LUN data from the snapshot. After completing the copy, the snapshot is deleted. The copy process may take several minutes to complete for large LUNs if the source LUN has a pre-allocated provisioning policy. The copy process may also take several minutes for thinly provisioned LUNs that are significantly utilized.
Shadow	Snapshot	Not allowed.
Shadow	Shadow	A snapshot of the shadow LUN is created. The clone LUN is configured to reference data from the snapshot.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
2. In the **Logical Units** section, under **Actions**, click **Clone** for the logical unit that you want to clone.
The **iSCSI Clone Logical Unit** page appears.
3. In the **LUN type** list, click **Normal**, **Shadow**, or **Snapshot**.
4. Modify the remaining settings as needed.
The configurable settings for the clone LUN vary according to the source LUN type.
5. Click **Submit**.

Move a logical unit

You can move a logical unit from one target to another, change the value of its logical unit number (LUN), or manually update the path to the LUN directory.



Note:

- The name of a logical unit is composed of its target name and LUN value, separated by a colon (for example, **mytarget:0**).
- You cannot modify the path of a snapshot LUN.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
2. In the **Logical Units** section, under **Actions**, click **Move** for the logical unit that you want to move.
The **iSCSI Move Logical Unit** page appears.
3. In the **To target** list, click to select a new target for the logical unit.
4. Click one of the **To LUN number** options.
 - To assign the next available number to the logical unit, click **Automatic**. This is the default setting.
 - To manually assign a number to the logical unit, click **Manual** and then, in the **Number** box, type an integer value. Valid values must be within the range 0-255, and must not be currently assigned to another logical unit.
5. To manually configure the path where the LUN directory will be created, in the **To path** box, type the full path of the directory, or click **Browse** to locate the directory.
The directory must be within the `/ifs` directory tree. The full path to the directory is required, and wildcard characters are not allowed. The default path is `/ifs/data/ISCSI.LUN.<TargetName>.<LUNnumber>`, where `<TargetName>` is defined in the **To target** field and `<LUNnumber>` is defined in the **To LUN number** field.
6. Click **Submit**.

Import a logical unit

You can recreate logical units that have been replicated to a remote cluster or that have been backed up and then restored to a remote cluster.

The iSCSI module does not support replicating or restoring logical unit snapshots or shadow copies.

1. On the **File System** menu, point to **iSCSI**, and then click **Targets & Logical Units**.
The **iSCSI Targets & Logical Units** page appears.
2. In the **Logical Units** section, click **Import logical unit**.
The **iSCSI Import Logical Unit** page appears.
3. Optionally, in the **Description** box, type a descriptive comment for the logical unit.
4. In the **Source path** box, type the full path (beginning with `/ifs`) of the directory that contains the logical unit that you want to import, or click **Browse** to locate the directory.
5. In the **Target** list, click to specify the target that will contain the logical unit.
6. Select one of the **LUN number** options.
 - To assign the next available number to the logical unit, click **Automatic**. This is the default setting.
 - To manually assign a number to the logical unit, click **Manual** and then, in the **Number** box, type an integer value. Valid values must be within the range 0-255, and must not be currently assigned to another logical unit.
7. Select one of the **LUN access** options.
 - To make the LUN accessible, click **Online**. This is the default setting.
 - To make the LUN inaccessible, click **Offline**.
8. Select one of the **Write access** options.
 - To allow iSCSI initiators to write to the LUN, click **Read-Write**. This is the default setting.
 - To prevent iSCSI initiators from writing to the LUN, click **Read-Only**.
9. Select one of the caching options.
 - To allow write caching for files storing LUN data, click **Enable**.
 - To prevent write caching for files storing LUN data, click **Disable**.
10. Click **Submit**.

Apache Hadoop

Apache Hadoop is a flexible, open-source framework for large-scale distributed computation. The OneFS file system can be configured for native support of the Hadoop Distributed File System (HDFS) protocol, enabling your cluster to participate in the Hadoop ecosystem.

HDFS integration requires a separate license. For additional information, or to activate HDFS support for your Isilon clustered storage system, contact your Isilon Systems sales representative.

OneFS supports the following HDFS distributions:

- Apache Hadoop 0.20.203
- Apache Hadoop 0.20.205
- Cloudera's distribution including Apache Hadoop version 3, update 1 (CDH3 Update 1)
- Greenplum HD 1.1

For more information about the installation, configuration, and use of Hadoop projects, visit <http://hadoop.apache.org>.

Isilon and Hadoop cluster integration

To enable native HDFS support in OneFS, you must integrate the Isilon cluster with a cluster of Hadoop compute nodes.

This process requires configuration of the Isilon cluster as well as each Hadoop compute node that needs access to the cluster.

Create a local Hadoop user

To enable access to files that are stored on OneFS via the HDFS protocol, you must first create a local Hadoop user that maps to a user on a Hadoop client.

This procedure describes how to create a Hadoop user through the OneFS web administration interface.

1. On the **File Sharing** menu, point to **Authentication Sources**, and then click **Local Users**.
The Authentication Sources **Local Users** page appears.
2. Click the **Add user** link.
The **Modify Local User** page appears.
3. Configure the following settings as shown. All other settings are optional.
 - **User name:** Type a user name that maps to a client-side Hadoop user.
 - **Home directory:** Type `/ifs/home/<user_name>`, where `<user_name>` matches the above **User name** setting.
 - **Shell:** Select the **zsh** shell.
4. Click **Submit**.

Configure a Hadoop client

To access HDFS on the Isilon cluster from a Hadoop client, you must modify the client's Hadoop configuration settings and start the MapReduce engine.



Note: This procedure performs the following actions on a Hadoop client node:

- Configures the client to use an HDFS-enabled Isilon cluster as the default file system.
- Disables speculative execution for MapReduce jobs.
- Starts the MapReduce engine, which consists of the JobTracker and TaskTracker services, in order to perform MapReduce jobs.

In the following steps, replace `$HADOOP_INSTALL` with the location of the Hadoop installation.

1. On the Hadoop client, navigate to the `$HADOOP_INSTALL/conf` directory.
2. Using a text editor such as `vi`, update the `core-site.xml` and `mapred-site.xml` configuration files to set the Isilon cluster as the default file system and to disable speculative execution.

- In the `core-site.xml` file, add or update the following lines, where [Node IP] is the IP address of any node in the Isilon cluster:

```
<property>
  <name>fs.default.name</name>
  <value>hdfs://[Node IP]:8020</value>
</property>
```

- In the `mapred-site.xml` file, add or update the following lines:

```
<property>
  <name>mapred.map.tasks.speculative.execution</name>
  <value>>false</value>
</property>
<property>
  <name>mapred.reduce.tasks.speculative.execution</name>
  <value>>false</value>
</property>
```

3. Navigate to the `$HADOOP_INSTALL/bin` directory.
4. Start the MapReduce engine by running the following script:


```
./start-mapred.sh
```

Configure HDFS

You can specify which HDFS distribution to use, the logging level, root path, and Hadoop block size, and the number of available worker threads.

HDFS configuration is performed by running the `isi hdfs` command in the OneFS command-line interface.

1. Open a secure shell (SSH) connection to any node in the cluster and log in using the root account.



Note: You can combine multiple options with a single `isi hdfs` command. For command usage and syntax, run the `isi hdfs -h` command.

2. To specify which distribution of the HDFS protocol to use, run the `isi hdfs` command with the `--force-version` option.

Valid values are listed below. Please note that these values are case-sensitive.

- **AUTO**: Attempts to match the distribution that is being used by the Hadoop compute node.
- **APACHE_0_20_203**: Uses the Apache Hadoop 0.20.203 release.
- **APACHE_0_20_205**: Uses the Apache Hadoop 0.20.205 release.
- **CLOUDERA_CDH3**: Uses version 3 of Cloudera's distribution including Apache Hadoop.
- **GREENPLUM_HD_1_1**: Uses the Greenplum HD 1.1 distribution.

For example, the following command forces OneFS to use version 0.20.203 of the Apache Hadoop distribution:

```
isi hdfs --force-version=APACHE_0_20_203
```

3. To set the default logging level for the Hadoop daemon across the cluster, run the `isi hdfs` command with the `--log-level` option.

Valid values are listed below, in descending order from highest to lowest logging level. The default value is **NOTICE**. Please note that these values are case-sensitive.

- **EMERG**: A panic condition. This is normally broadcast to all users.
- **ALERT**: A condition that should be corrected immediately, such as a corrupted system database.
- **CRIT**: Critical conditions, such as hard device errors.
- **ERR**: Errors.
- **WARNING**: Warning messages.
- **NOTICE**: Conditions that are not error conditions, but may need special handling.
- **INFO**: Informational messages.
- **DEBUG**: Messages that contain information normally of use only when debugging a program.

For example, the following command sets the log level to **WARNING**:

```
isi hdfs --log-level=WARNING
```

4. To set the path on the cluster to present as the HDFS root directory, run the `isi hdfs` command with the `--root-path` option.

Valid values include any directory path beginning at `/ifs`, which is the default HDFS root directory.

For example, the following command sets the root path to `/ifs/hadoop`:

```
isi hdfs --root-path=/ifs/hadoop
```

5. To set the Hadoop block size, run the `isi hdfs` command with the `--block-size` option.

Valid values are **4KB** to **1GB**. The default value is **64MB**.

For example, the following command sets the block size to 32 MB:

```
isi hdfs --block-size=32MB
```

6. To tune the number of worker threads that HDFS uses, run the `isi hdfs` command with the `--num-threads` option.

Valid values are **1** to **256** or **auto**, which is calculated as twice the number of cores. The default value is **auto**.

For example, the following command specifies 8 worker threads:

```
isi hdfs --num-threads=8
```

7. To allocate IP addresses from an IP address pool, run `isi hdfs` with the `--ip-pool` option.

Valid values are in the form `<subnet>:<pool>`.

For example, the following command allocates IP addresses from a pool named 'pool2', which is in the 'subnet0' subnet:

```
isi hdfs --ip-pool=subnet0:pool2
```

Enable or disable the HDFS service

The HDFS service, which is enabled by default, can be manually enabled or disabled by running the `isi services` command.

1. Open a secure shell (SSH) connection to any node in the cluster and log in using the root user account.
2. At the command prompt, run the `isi service` command to enable or disable the HDFS service, `isi_hdfs_d`.
 - To enable the HDFS service, run:
`isi service isi_hdfs_d enable`
 - To disable the HDFS service, run:
`isi service isi_hdfs_d disable`

The Isilon SupportIQ module

The Isilon SupportIQ module is a tool that, when enabled, allows Isilon Technical Support personnel to request scripts that gather diagnostic data about your cluster, and then upload the data securely to Isilon. The data collected by SupportIQ can assist Isilon Technical Support personnel in troubleshooting cluster issues without the need for you to manually execute and upload log files, and can shorten the time required to resolve issues.

The SupportIQ module is included with the OneFS operating system, and does not require a separate license. SupportIQ is enabled by default; however, you can disable, enable, and configure SupportIQ through the Isilon web administration interface.

SupportIQ can run two different types of scripts:

- **Scripts that gather data:** These scripts can be run automatically, at the request of an Isilon Technical Support representative, and collect information about your cluster's configuration settings and operations. The SupportIQ agent then uploads the gathered information to a secure Isilon FTP site so it is available for Isilon Technical Support personnel to review. These scripts do not affect any cluster services or the availability of your data. These scripts are based on the Isilon `isi_gather_info` log-gathering tool; for more information, see the `isi_gather_info` man page.
- **Scripts that can affect services:** The SupportIQ agent runs these scripts on demand, as initiated remotely by authorized Isilon Technical Support personnel. These scripts collect and upload information about your cluster's configuration settings and operations, and can affect cluster services. Before initiating a service-affecting script, an Isilon Technical Support representative will contact you to confirm the impending script run.

Remote access

In addition to enabling the SupportIQ module, which allows the SupportIQ agent to run scripts, you can also optionally enable remote access, which allows Isilon Technical Support personnel to remotely manage your cluster through the SSH-based command-line interface or the web administration interface. Remote access can enable Isilon Technical Support to more quickly identify and troubleshoot cluster issues. If you enable remote access, you must also share your cluster login credentials with the appropriate Isilon Technical Support personnel. Isilon Technical Support personnel will remotely access your cluster only in the context of an open Support case, and only after receiving your explicit permission to remotely connect to your cluster.

Data collected by SupportIQ

The Isilon SupportIQ module collects basic data about your cluster, including usage statistics, alerts related to cluster events, and information about hardware configuration, OneFS versions, and installed software patches. SupportIQ does not collect any information related to the data stored on your cluster.

SupportIQ collects the following data:

- Cluster capacity and usage (total capacity, amount used, and percent used)
- Group management protocol (GMP) status
- Cluster GUID (unique cluster identifier)
- Installed patches for Isilon software products
- Active Directory join status
- LDAP join status
- Cluster name
- Number of nodes in cluster

- Software licenses
- Time zone
- Count of active SMB connections to a specified node
- Count of active FTP connections to a specified node
- Count of active HTTP connections to a specified node
- Count of active NFS connections to a specified node
- LSI firmware version
- Mellanox firmware version
- Vitesse firmware version
- InfiniBand firmware version
- EX-node configuration status
- Node memory (installed RAM)
- Node hardware family
- Node model name
- Node hardware part number
- Node ID
- Logical node number
- Load average
- CPU utilization percentage
- Node serial number
- Node status (health indicator)
- Time since last reboot
- Dual InfiniBand status (enabled or disabled)
- InfiniBand subnet master status (enabled or disabled)
- OneFS version
- `/var/crash` partition capacity and usage (total capacity, amount used, and percent used)
- `/ifs` partition capacity and usage (total capacity, amount used, and percent used)
- `/` partition capacity and usage (total capacity, amount used, and percent used)
- `/var` partition capacity and usage (total capacity, amount used, and percent used)
- Leak freed blocks syscontrol status (on or off)
- SupportIQ agent version
- SupportIQ module status (enabled or disabled)
- SupportIQ remote-access status (enabled or disabled)

SupportIQ scripts

If SupportIQ is enabled, Isilon Technical Support personnel can request logs via scripts that gather cluster data and then upload the data to Isilon or, with your approval, execute a predefined set of commands on the cluster.

All SupportIQ scripts are located in the `/usr/local/SupportIQ/Scripts/` directory on each node.

Data-gathering scripts

The following table lists all of the data-gathering scripts that SupportIQ can run if SupportIQ is enabled. These scripts can be run automatically, at the request of an Isilon Technical Support representative, and collect information about your cluster's configuration settings and operations. The SupportIQ agent then uploads the gathered information to a secure Isilon FTP site so it is available for Isilon Technical Support personnel to review. These scripts do not affect any cluster services or the availability of your data. These scripts are based on the Isilon `isi_gather_info` log-gathering tool; for more information, see the `isi_gather_info` man page.

Script name	Description
GenerateDashboardFile.sh	Generates and uploads the dashboard.xml file. This script invokes the isi_check_cluster command, which generates all of the SupportIQ data elements.
GetData.sh	Master data-collection script; invoked by other scripts to gather specific data items.
GetData-application.sh	Collects and uploads information about OneFS application programs, and invokes the GetData.sh script.
GetData-cluster.sh	Collects and uploads information about overall cluster configuration and operations, and invokes the GetData.sh script.
GetData-domain.sh	Collects and uploads information about the cluster's Active Directory Services (ADS) domain membership, and invokes the GetData.sh script.
GetData-fs.sh	Collects and uploads information about the state and health of the OneFS /ifs/ file system, and invokes the GetData.sh script.
GetData-ib.sh	Collects and uploads information about the configuration and operation of the InfiniBand back-end network, and invokes the GetData.sh script.
GetData-logs.sh	Initiates the isi_gather_info command, which collects and uploads only the most recent cluster log information, and invokes the GetData.sh script.
GetData-messages.sh	Collects and uploads any currently active /var/log/messages files, and invokes the GetData.sh script.
GetData-network.sh	Collects and uploads information about cluster-wide and node-specific network configuration settings and operations, and invokes the GetData.sh script.
GetData-node.sh	Collects and uploads node-specific configuration, status, and operational information, and invokes the GetData.sh script.
GetData-protocol.sh	Collects and uploads network status information and configuration settings for the NFS, SMB, FTP, and HTTP protocols, and invokes the GetData.sh script.
GetData-usage.sh	Collects and uploads current and historical information about node performance and resource usage, and invokes the GetData.sh script.
IsiGatherInfo.sh	Initiates the isi_gather_info command, which collects and uploads recent cluster log information.
IsiGatherInfo-incremental.sh	Initiates an incremental (since the most recent full isi_gather_info operation) version of the isi_gather_info command, which collects and uploads recent cluster log information.
UploadDashboardFile.sh	Uploads a copy of the dashboard.xml file to the secure Isilon Technical Support FTP site.

Service-affecting scripts

The following table lists all of the service-affecting scripts that SupportIQ can run if SupportIQ is enabled. These scripts collect and upload information about your cluster's configuration settings and operations, and can affect cluster services. Only authorized senior-level Isilon Technical Support engineers can run these scripts, and only after receiving your explicit permission.

Script name	Description
CleanVarCrash.sh	Deletes any files that were previously uploaded by SupportIQ.
RebootNode.sh	Reboots a specific node.
RestartAgent.sh	Restarts the SupportIQ agent, and invokes the RestartDaemon.sh script.
RestartDaemon.sh	Master daemon-restart script; invoked by other scripts to restart specific daemons.
RestartDaemon-lsassd.sh	Restarts the lsassd authentication-services daemon on all nodes, and invokes the RestartDaemon.sh script.
RestartService.sh	Master service-restart script; invoked by other scripts to restart specific services.
RestartService-apache2.sh	Restarts the OneFS Apache2 service, and invokes the RestartService.sh script.
RestartService-isi_lcd.sh	Restarts the OneFS isi_lcd service, and invokes the RestartService.sh script.
RestartService-isi_webui.sh	Restarts the OneFS isi_webui service, and invokes the RestartService.sh script.
RestartService-nfs.sh	Restarts the OneFS nfs service, and invokes the RestartService.sh script.
RestartService-samba.sh	Restarts the OneFS samba service, and invokes the RestartService.sh script.

SupportIQ configuration and management

You must enable and configure the SupportIQ module before SupportIQ can run any scripts and start gathering data. The SupportIQ module is enabled by default.

You can also optionally enable remote access, which allows Isilon Technical Support personnel to remotely manage your cluster. Remote access is disabled by default. To enable Isilon to remotely access your cluster, you must provide the cluster password to the appropriate Technical Support engineer.

Enable and configure SupportIQ

If you enable the Isilon SupportIQ module, the SupportIQ agent can run scripts that gather and upload data. You can also optionally enable remote access to your cluster.

1. On the **Cluster** menu, point to **Cluster Settings**, and then click **SupportIQ**.
The **SupportIQ** page appears.
2. In the **SupportIQ Settings** area, select the **Enable SupportIQ** check box.
3. Click to specify a **SupportIQ alerts** option:
 - **Send alerts via SupportIQ agent (HTTPS) and by email (SMTP):** SupportIQ delivers alerts to Isilon through the SupportIQ agent over HTTPS, and by email over SMTP.
 - **Send alerts via SupportIQ agent (HTTPS):** SupportIQ delivers alerts to Isilon only through the SupportIQ agent over HTTPS.
4. Optionally, enable HTTP proxy support for SupportIQ:
 - a. Select the **Enable HTTP proxy for SupportIQ** check box.
 - b. In the **Proxy host** box, type the IP address or fully qualified domain name (FQDN) of the HTTP proxy server.
This field is required.

- c. In the **Proxy port** box, type the number of the port on which the HTTP proxy server listens for requests.
This field is required.
- d. In the **Username** box, type the user name for the proxy server, if applicable.
- e. In the **Password** box, type the password for the proxy server, if applicable.
5. Optionally, enable remote access to the cluster:
 - a. Select the **Enable remote access to cluster via SSH and web interface** check box.
The remote-access end user license agreement (EULA) appears.
 - b. Review the EULA and, if you agree to the terms and conditions, select the **I have read and agree to...** check box.
6. Click **Submit**.

Disable SupportIQ

If you disable SupportIQ, the SupportIQ agent can no longer run scripts that gather and upload data, and your cluster data is no longer visible to Isilon Technical Support personnel.

1. On the **Cluster** menu, point to **Cluster Settings**, and then click **SupportIQ**.
The **SupportIQ** page appears.
2. In the **SupportIQ Settings** area, clear the **Enable SupportIQ** check box.
3. Click **Submit**.

Events and event notifications

Events and event notifications enable you to receive information about the health and performance of the cluster, including drives, nodes, snapshots, network traffic, and hardware.

As part of working with events, you can:

- View information about specific events.
- Create and configure event notifications.
- Troubleshoot specific events.

View events and event notifications

You can use the **Events > Summary** and **Events > History** pages to view events, notification rules, events settings, and event-specific help.

View the event summary

The **Events > Summary** page provides an overview of all currently generated events on a cluster. You can use it to view and manage new events, open events, and recently ended events. You can also view coalesced events, and navigate to detailed event pages to view additional information about specific events.

Select the **Status** menu, point to **Events**, and click **Summary**.

The **Events > Summary** page appears and displays the following sections:

- **Events:** Shows a list of all recently generated events. Select **Quiet all events** to quiet the events that appear in this list.
- **Quieted Events:** Shows all recently quieted events.
- **Ended Events:** Shows a list of recently ended events. Select **View all events** to view the **Events > History** page.

View the event history

The **Events > History** page provides a list of previous events that no longer appear on the **Events > Summary** page. You can also view older coalesced events, and navigate to detailed event pages to view additional information about specific events.

Select the **Status** menu, point to **Events**, and click **History**.

The **Events > History** page appears and displays the following sections:

- **Event History:** Shows the event history list. You can select **Quiet all events** to quiet the events in the list and select **Show coalesced events** to show or hide any coalesced events in the event history.

View event notification rules

The **Events > Notification Rules** page provides a list of all notification rules created on a cluster. You can view and edit existing notification rules, and add new rules.

Select the **Status** menu, point to **Events**, and click **Notification Rules**.

The **Events > Notification Rules** page appears and displays the following sections:

- **Notification Rules:** Shows all notification rules created on a cluster. Select **Add rule** to add a new notification rule.

View events settings

The **Events > Settings** page allows you to modify contact information, email settings, and SupportIQ settings on a cluster. You can also send a test event from this page.

Select the **Status** menu, point to **Events**, and click **Settings**.

The **Events > Settings** page appears and displays the following sections:

- **Send Test Event:** Shows a link to generate a test event.
- **SupportIQ:** Shows the current SupportIQ settings. To change the current settings, select **Modify SupportIQ settings**.
- **Email Settings:** Shows the current email settings for the cluster. To change the current settings, select **Modify Email settings**.
- **Contact Information:** Shows the current contact information settings. To change the current settings, select **Modify Contact Information settings**.

View event help

You can view event help for all generated events that appear on the **Event > Summary** and **Event > History** pages.

Select **View** in the **Actions** column of an event list.

The **Events > View Event** page appears and displays the following sections:

- **Event Details:** Shows specific information about the event, such as start time, event message, and scope. Select **View Event Help** to view additional information about the event and any actions you can take to resolve issues related to the event.
- **Coalesced Events:** If the event is a coalescing event, this section is displayed and shows all events related to the coalescing event.

Event management

You can view, quiet, and cancel events using the **Events > Event Summary** page. Additionally, you can send a test event to verify that your contact information allows you to receive notifications.

Quiet an event

Perform the following steps to quiet an event.

Select **Quiet** in the **Actions** column of an event list.

The event no longer appears in the **Event > Summary** list.

Send a test event

Perform the following steps to generate a test event.

1. Select the **Status** menu, point to **Events**, and click **Settings**.
The **Events > Settings** page appears.
2. Select **Send Test Event**.
A test event is generated. This allows you to verify your contact information so you can properly receive event notifications.

Event notification management

You can add and edit event notifications through the **Events > Event Notifications** page.

Add an event notification rule

The **Events > Notification Rules** page allows you to add a new notification rule.

1. Select the **Status** menu, point to **Events**, and click **Notification Rules**.
The **Events > Notification Rules** page appears.
2. Select **Add rule**.
The **Events > Add Notification Rule** page appears.

Edit an event notification rule

The **Events > Notification Rules** page allows you to edit an event notification rule.

1. Select the **Status** menu, point to **Events**, and click **Notification Rules**.
The **Events > Notification Rules** page appears.
2. Select **Edit** in the **Actions** column for the notification rule.
The **Events > Edit Notification Rule** page appears.

Event notification reference

System disk events

System disk events provide information about system disk status, such as drive availability, disk errors, and partition usage.

100010001

The `/var` partition is near capacity.

Details

The `/var` partition on the node is at or near its capacity.

Administrator Response

If, after you view this event notification, the `/var` partition returns to a normal usage level of less than 75 percent, and the issue does not recur, you can disregard the event.

View the `df` output for the `/var` partition, which is included in the event text. Depending on the output, perform one or more of the following tasks:

- If the `/var` partition returns to a normal usage level, review the list of recently written logs to determine whether a specific log is rotating frequently.

Rotation can resolve the full-partition issue by compressing or removing large logs and old logs, thereby automatically reducing partition usage.
- If the Samba log is rotating frequently, check the Server Message Block (SMB) debug levels. If the debug levels are set to a value greater than 1 and the system is not running in debug mode, reset the debug levels to 1.

- If other logs are rotating frequently, or if the preceding tasks do not resolve the issue, run the `isi_gather_info` command to gather logs, and then contact Isilon Technical Support.

100010002

The `/var/crash` partition is near capacity.

Details

Data is filling the 512 MB `/var/crash` partition on a node. The full text of the event includes a percentage-used amount on the partition.

The purpose of the partition is to preserve data about failed processes and unplanned restarts to enable analysis of those events.

Administrator Response

If you are temporarily storing files in the `/var/crash` partition, remove them. If you need to keep the files in the `/var/crash` partition for more time, quiet the event.

If you are not temporarily storing files in the `/var/crash` partition, contact Isilon Technical Support.

100010003

Capacity Used - Root Partition

Details

The root partition on the reporting node has exceeded its specified space limit.

A common cause of this condition is that files were copied to the root user's home directory.

Isilon Systems does not support storing user data on the root partition of a node. User data that was accidentally copied to the root partition via secure copy or an alternate method may cause a wide variety of problems.

Administrator Response

Delete data from the root partition or move it to the `/ifs` partition.

100010005

A SAS PHY topology problem or change was detected.

Details

The SAS PHY monitor detected an error or a change in the disk subsystem.

If the drive was recently replaced, you can safely ignore this event.

Administrator Response

Ensure that the reported drive and cable are securely seated.

If the issue is not resolved, contact Isilon Technical Support.

100010006

A drive's error log counter indicates there may be a problem.

Details

The drive's error log detected an error or a change in the disk subsystem.

Administrator Response

Ensure that the reported drive and cable are securely seated.

If the issue is not resolved, contact Isilon Technical Support.

100010007

The SAS link connected to the chassis has exceeded the maximum bit error rate.

Details

The SAS link connected to the chassis has exceeded the maximum bit error rate.

Administrator Response

Ensure that the reported drive and cable are securely seated.

If the issue is not resolved, contact Isilon Technical Support.

100010008

The SAS link connected to the chassis has been disabled for exceeding the maximum bit error rate.

Details

The SAS link connected to the chassis has been disabled for exceeding the maximum bit error rate.

Administrator Response

Ensure that the reported drive and cable are securely seated.

If the issue is not resolved, contact Isilon Technical Support.

100010009

Node has drive in bay in restripe.

Details

Node has drive in bay in restripe.

Administrator Response

No action is required. The message is informational only.

100010010

Node drive in bay removed.

Details

One or more drive bays on a node have a drive that has been removed.

Administrator Response

Complete the following steps:

1. Replace the drive.
2. Optional: To quiet this event, run `isi_healthy_drives_alert disable`.

100010011

One or more drives are ready to be replaced.

Details

One or more failed drives in this node are ready to be replaced.

Administrator Response

Complete the following steps:

1. In `/var/logs/messages`, review the messages that were generated during the most recent restriper operation, and then perform one of the following steps:
 - If the restriper operation completed successfully, replace the drive.
 - If the restriper operation failed, contact Isilon Technical Support.
2. Optionally, disable the event by running the following command:
`isi_healthy_drives_alert disable`

100010013

ECC retry.

Details

ECC retry.

Administrator Response

No action is required. The message is informational only.

100010014

ECC list is full.

Details

ECC list is full.

Administrator Response

No action is required. The message is informational only.

100010015

The SmartPool is near capacity.

Details

A SmartPool is reaching or has reached full capacity.

Administrator Response

No action is required. The message is informational only.

100010016

The SmartPool did not satisfy SSD layout preference

Details

A SmartPool did not satisfy the SSD layout preference.

Administrator Response

The message is informational only.

199990001

Disk errors detected.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

Node status events

Node status events provide information about system disk status, such as node availability and performance.

200010001

Node offline.

Details

One of the following conditions is true:

- One or more nodes are offline.
- A node lacks back-end network connectivity. (Back-end connectivity refers to a node's ability to communicate with other nodes.)
- A node cannot join the group.

Administrator Response

Perform the following steps:

1. Determine whether the node is turned on. Visually inspect the node to verify that the power light is on, or check the node via an intelligent power management device.
2. If the node is turned off, attempt to turn it on.
 - If the node turns on, monitor it to determine whether it joins the cluster.
 - If the node rejoins the cluster, run the `isi_gather_info` command, and then contact Isilon Technical Support for failure analysis.
 - If the node does not rejoin the cluster, proceed to step 3.
 - If the node does not turn on, ensure that the circuit breakers are operational and that the power outlets are active.

- If the node is receiving power, contact Isilon Technical Support for help troubleshooting the issue.
 - If the node is not receiving power, resolve the power supply issue.
3. If the node is on but did not rejoin the cluster, attempt to establish remote access via a secure shell (SSH) session. If the SSH session fails, attempt to establish remote access via the serial console.
 4. If neither the SSH session nor the serial console is responsive, press CTRL+T either in the SSH session or on the serial console.
 - If pressing CTRL+T produces output, record the output, and then contact Isilon Technical Support for failure analysis.
 - If the node is unresponsive, turn the power off and then on again. Run the `isi_gather_info` command, and then contact Isilon Technical support for failure analysis.
 5. Contact Isilon Technical Support for further software troubleshooting.

200010002

Node online.

Details

A node that was previously offline has rejoined the group.

Administrator Response

No action is required. The message is informational only.

200010003

Node has been offline for over a week.

Details

One of the following conditions is true:

- One or more nodes are offline.
- A node lacks back-end network connectivity. Back-end connectivity refers to a node's ability to communicate with other nodes.
- A node cannot join the group.

Administrator Response

Perform the following steps:

1. Determine whether the node is turned on: Visually inspect the node to verify that the power light is on, or check the node via an intelligent power management device.
2. If the node is turned off, attempt to turn it on.
 - If the node turns on, monitor it to determine whether it joins the cluster.
 - If the node rejoins the cluster, run the `isi_gather_info` command, and then contact Isilon Technical Support for failure analysis.
 - If the node does not rejoin the cluster, proceed to step 3.
 - If the node does not turn on, ensure that the circuit breakers are operational and that the power outlets are active.
 - If the node is receiving power, contact Isilon Technical Support for help with troubleshooting the issue.
 - If the node is not receiving power, resolve the power supply issue.

3. If the node is on but did not rejoin the cluster, attempt to establish remote access via a secure shell (SSH) session. If the SSH session fails, attempt to establish remote access via the serial console.
4. If neither the SSH session nor the serial console is responsive, press CTRL+T either in the SSH session or on the serial console.
 - If pressing CTRL+T produces output, record the output, and then contact Isilon Technical Support for failure analysis.
 - If the node is unresponsive, turn the power off and then on again. Run the `isi_gather_info` command, and then contact Isilon Technical support for failure analysis.
5. Contact Isilon Technical Support for further troubleshooting assistance.

200010004

Elevated CPU usage

Details

A node in the cluster reports elevated CPU usage for an extended period of time.

Administrator Response

The appropriate action to take depends upon the circumstances:

- If the event appeared shortly after a restriper operation began, it is recommended that you wait until the restriper operation is complete before investigating the cause.
- If the cluster is experiencing high throughput, high CPU usage is expected and you can safely suppress this event via the Isilon web administration interface.

High cluster traffic causes the `smbd` process, the `nfsd` process, or both to use a large number of CPU cycles.

If neither of the above circumstances is responsible for the high CPU usage, run the `isi_gather_info` command, and then contact Isilon Technical Support for help investigating the cause.

200010005

EX (Disk-Only) Node Offline

Details

An Isilon EX node that is connected to another node is turned off or otherwise unavailable. An EX node is sometimes referred to as a "disk-only node."

If maintenance was recently performed in the data center, it is possible that the node or cable was inadvertently bumped out of position.

Administrator Response

Attempt the following possible solutions in the order presented. If a solution resolves the issue, there is no need to attempt the other possible solutions.

- Ensure that the cable to the node is firmly seated.
- Determine whether the EX node is turned on.
 - If it is turned off, attempt to turn it back on. Note whether the node turns on, and then contact Isilon Technical Support.
 - If it is turned on, turn off the head node. (Each EX node is connected to one head node.) When the head node is turned off, turn the EX node off and then on again, and then turn on the head node. When the EX node restarts, note whether the head node detects the EX node, and then contact Isilon Technical Support.

200020001

Gigabit Ethernet link running below capacity.

Details

An Ethernet link is not operating at maximum throughput.

Administrator Response

Attempt the following possible solutions in the order that they are presented. If a solution resolves the issue, there is no need to attempt other possible solutions.

- If this event persists but throughput operation is not a recurring issue, verify the following items:
 - The node's network cable is connected securely.
 - The cable is rated for the appropriate Ethernet speed.
 - The switch port speed is set to the same (or higher) speed as the Ethernet cable.
 - The switch port is set to the Full Duplex setting.
- If the cable is connected securely, plug the cable into a different node that has a network port functioning at full speed and that has an identical network configuration. (When you plug the cable into the other node, leave the other end of the cable plugged into the same switch port.)
- If the issue persists, replace the cable.
- If the issue persists, move the cable to another port on the switch.
- If the issue persists, review the switch logs and consult your switch user manual.
- If this event is not ongoing, and this is not a recurring problem, check for recent maintenance in the rack or adjoining racks as a cable may have been disconnected during maintenance.

If none of these possible solutions resolves the issue, contact Isilon Technical Support for further troubleshooting assistance.

200020002

10-Gigabit Ethernet link running below capacity.

Details

A 10-Gigabit Ethernet link is not operating at maximum throughput.

Administrator Response

Attempt the following possible solutions in the order that they are presented. If a solution resolves the issue, there is no need to attempt other possible solutions.

- If this alert persists but throughput operation is not a recurring issue, verify the following items:
 - The node's network cable is connected securely.
 - The cable is rated for the appropriate Ethernet speed.
 - The switch port speed is set to the same (or higher) speed as the Ethernet cable.
 - The switch port is set to the Full Duplex setting.
- If the cable is connected securely, plug the cable into a different node that has a network port functioning at full speed and that has an identical network configuration. (When you plug the cable into the other node, leave the other end of the cable plugged into the same switch port.)
- If the issue persists, replace the cable.
- If the issue persists, move the cable to another port on the switch.
- If the issue persists, review the switch logs and consult your switch user manual.

- If this alert is not ongoing, and this is not a recurring problem, check for recent maintenance in the rack or adjoining racks as a cable may have been disconnected during maintenance.

If none of these possible solutions resolves the issue, contact Isilon Technical Support for further troubleshooting assistance.

200020003

Internal network link down.

Details

One or more interfaces on a node have lost carrier.

Administrator Response

Attempt the following possible solutions in the order presented. If a solution resolves the issue, there is no need to attempt the other possible solutions.

- Ensure that the node's network cable is connected securely and that neither the cable nor the connector is damaged.
- If the carrier issue is ongoing but does not continually recur, attempt the following possible solutions:
 - Plug the cable into another node that has a functioning network port and an identical network configuration. Be sure to plug the cable into the same switch port on the different node. Determine whether the issue is related to the cable or the node.
 - If the issue persists, replace the cable.
 - If the issue still persists, move the cable to a different port on the switch.
 - If the issue still persists, review the switch logs.
- If this event is not ongoing and this is not a recurring problem, ensure that all cables in the rack and in adjacent racks are connected securely.

If none of these possible solutions resolves the issue, contact Isilon Technical Support for troubleshooting assistance.

200020004

Aggregate network link error

Details

Aggregate network link error.

Administrator Response

Attempt the following possible solutions in the order that they are presented. If a solution resolves the issue, there is no need to attempt other possible solutions.

- If this alert persists but throughput operation is not a recurring issue, verify the following items:
 - The node's network cable is connected securely.
 - The cable is rated for the appropriate Ethernet speed.
 - The switch port speed is set to the same (or higher) speed as the Ethernet cable.
 - The switch port is set to the Full Duplex setting.
- If the cable is connected securely, plug the cable into a different node that has a network port functioning at full speed and that has an identical network configuration. (When you plug the cable into the other node, leave the other end of the cable plugged into the same switch port.)
- If the issue persists, replace the cable.
- If the issue persists, move the cable to another port on the switch.
- If the issue persists, review the switch logs and consult your switch user manual.

- If this alert is not ongoing, and this is not a recurring problem, check for recent maintenance in the rack or adjoining racks as a cable may have been disconnected during maintenance.

If none of these possible solutions resolves the issue, contact Isilon Technical Support for further troubleshooting assistance.

200020005

External network link down.

Details

One or more interfaces on a node have lost carrier.

Administrator Response

Attempt the following possible solutions in the order presented. If a solution resolves the issue, there is no need to attempt the other possible solutions.

- Ensure that the node's network cable is connected securely and that neither the cable nor the connector is damaged.
- If the carrier issue is ongoing but does not continually recur, attempt the following possible solutions:
 - Plug the cable into another node that has a functioning network port and an identical network configuration. Be sure to plug the cable into the same switch port on the different node. Determine whether the issue is related to the cable or the node.
 - If the issue persists, replace the cable.
 - If the issue still persists, move the cable to a different port on the switch.
 - If the issue still persists, review the switch logs.
- If this event is not ongoing and this is not a recurring problem, ensure that all cables in the rack and in adjacent racks are connected securely.

If none of these possible solutions resolves the issue, contact Isilon Technical Support for troubleshooting assistance.

299990001

Node status.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

299990002

Multiple internal network problems detected.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

299990003

Multiple external network problems detected.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

Reboot events

Reboot events provide information about node status, such as reboot operations and errors.

300010001

Node being rebooted.

Details

Node being rebooted.

Administrator Response

No action is required. The message is informational only.

300010002

Node being shut down.

Details

The node is being shut down.

Administrator Response

No action is required. The message is informational only.

300010003

Node reboot timed out.

Details

One of the following conditions is true:

- One or more nodes are offline.
- A node lacks back-end network connectivity. (Back-end connectivity refers to a node's ability to communicate with other nodes.)

- A node cannot join the group.

Administrator Response

Perform the following steps:

1. Determine whether the node is turned on. Visually inspect the node to verify that the power light is on, or check the node via an intelligent power management device.
2. If the node is turned off, attempt to turn it on.
 - If the node turns on, monitor it to determine whether it joins the cluster.
 - If the node rejoins the cluster, run the `isi_gather_info` command, and then contact Isilon Technical Support for failure analysis.
 - If the node does not rejoin the cluster, proceed to step 3.
 - If the node does not turn on, ensure that the circuit breakers are operational and that the power outlets are active.
 - If the node is receiving power, contact Isilon Technical Support for help troubleshooting the issue.
 - If the node is not receiving power, resolve the power supply issue.
3. If the node is on but did not rejoin the cluster, attempt to establish remote access via a secure shell (SSH) session. If the SSH session fails, attempt to establish remote access via the serial console.
4. If neither the SSH session nor the serial console is responsive, press CTRL+T either in the SSH session or on the serial console.
 - If pressing CTRL+T produces output, record the output, and then contact Isilon Technical Support for failure analysis.
 - If the node is unresponsive, turn the power off and then on again. Run the `isi_gather_info` command, and then contact Isilon Technical support for failure analysis.
5. Contact Isilon Technical Support for further software troubleshooting.

300020001

Node read-only transition failed.

Details

Node read-only transition failed.

Administrator Response

No action is required. The message is informational only.

300020002

Node journal backup validation failed.

Details

Node journal backup validation failed.

Administrator Response

No action is required. The message is informational only.

300020003

Node encountered an error performing final shutdown.

Details

Node encountered an error performing final shutdown.

Administrator Response

No action is required. The message is informational only.

399990001

Node is being rebooted.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

Software events

Software events provide information about OneFS and related application software status, such as SyncIQ policy issues and errors.

400030001

Process failed.

Details

A process failed to restart, despite several attempts to start it.

Administrator Response

Run `isi_gather_info`, and then contact Isilon Technical Support.

400030002

Process killed.

400040002

SyncIQ Policy Failure

Details

A policy failed to run.

Administrator Response

The appropriate action depends on the specific text of the event. Follow instructions as contained in the text of the event.

If you cannot resolve the issue, run `isi_gather_info` to gather logs, and then contact Isilon Technical Support.

400040005

SyncIQ unreplicated files in policy.

Details

An Isilon SyncIQ job failed to replicate a file to the target cluster.

The problem occurred because the file was deleted after it was first scanned, but before it was copied to the target cluster.

Administrator Response

Contact Isilon Technical Support.

400040008

Sync IQ Error

Details

A SyncIQ worker on the target cluster failed.

Generally, this error occurs as a side effect of another other problem that caused the target cluster to be unwriteable or a node to be unresponsive.

Administrator Response

Troubleshoot the issue as you would for an unwriteable cluster or an unresponsive node.

If you cannot resolve the issue, contact Isilon Technical Support.

400050001

Test event.

Details

A user requested that a test event be generated, either via the Isilon web administration interface or by running the `isi alert test` command.

Administrator Response

No action is required. The message is informational only.

400060001

AVScan No ICAP Servers Configured.

Details

The AVScan service is enabled, but no virus-scanning server has been configured.

Administrator Response

Perform the following steps:

1. Perform either of the following tasks:
 - If you do not intend to configure a supported Internet Content Adaptation Protocol (ICAP) antivirus server, disable the antivirus service via the Isilon web administration interface.
 - Configure an ICAP antivirus server in the web administration interface.

2. Cancel the existing event.
3. If the event persists, contact Isilon Technical Support.

400060002

AVScan No ICAP Servers Available.

Details

The Isilon IQ cluster cannot reach any antivirus server.

Administrator Response

Check the status of the reported antivirus servers and their network connectivity.

If you need assistance, contact Isilon Technical Support.

400060003

AVScan ICAP Server Unresponsive

Details

One antivirus server cannot be reached.

Administrator Response

Check network connectivity for the antivirus server for which the event was generated.

If you cannot resolve the issue, contact Isilon Technical Support.

400060004

AVScan Infected File Found.

Details

The virus-scanning software has identified a file that is infected by a virus.

Administrator Response

Review the antivirus scan report and address quarantined files as appropriate.

400070004

One or more licenses will expire soon.

Details

An evaluation license for a OneFS software module is scheduled to expire soon.

Administrator Response

To purchase the software module before the evaluation license expires, contact your Isilon sales representative.

400070005

One or more licenses have expired.

Details

An evaluation license for a OneFS software module has expired.

Administrator Response

To purchase the software module, contact your Isilon sales representative.

400080001

A firmware update has not been completely applied.

Details

A firmware update has either failed or not been applied.

Administrator Response

Attempt to reapply the firmware. From the node that reported the error, run the following command:

```
isi firmware update --local
```

400090001

Monthly Status.

Details

This alert is generated once each month to provide general cluster information.

Administrator Response

No action is required. The message is informational only.

400100001

Job state.

Details

Job state.

Administrator Response

No action is required. The message is informational only.

400100002

Job phase begin.

Details

Job phase begin.

Administrator Response

No action is required. The message is informational only.

400100003

Job phase end.

Details

Job phase end.

Administrator Response

No action is required. The message is informational only.

400100004

Job has failed.

Details

Job has failed.

Administrator Response

No action is required. The message is informational only.

400100005

Job policy

Details

Job policy

Administrator Response

No action is required. The message is informational only.

400110001

PID was killed to free pages.

Details

PID was killed to free pages.

Administrator Response

No action is required. The message is informational only.

499940001

SyncIQ is encountering problems with a policy.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

499960001

AVscan is encountering problems.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

SmartQuotas events

SmartQuotas events provide information about SmartQuotas status, such as report generation and configuration errors.

500010001

SmartQuotas

Details

The Isilon SmartQuotas module has notified a user of a quota violation.

Administrator Response

Action is optional. This event is information only.

If you do not want to receive events when SmartQuotas notifies a user of a quota violation, disable the event via the Isilon web administration interface.

If you have difficulty disabling the event, or if you believe that the system generated the event in error, contact Isilon Technical Support.

500010002

SmartQuotas notification for user failed.

Details

This error typically occurs as a result of one or more of the following conditions:

- The mail server is configured incorrectly.
- The mail server or the authentication server is down.
- A quota's address mapping rule is configured incorrectly.

Administrator Response

Review the quota configuration settings, and correct any apparent errors.

If you cannot resolve the issue, run the `isi_gather_info` command, gather the system output, and then contact Isilon Technical Support.

500010003

A SmartQuotas configuration error occurred in the config file.

Details

The Isilon SmartQuotas configuration is corrupt or invalid.

Administrator Response

Run the `isi_gather_info` command, and then contact Isilon Technical Support.

500010004

A SmartQuotas internal error occurred.

Details

The Isilon SmartQuotas module encountered an unexpected error.

Administrator Response

Run the `isi_gather_info` command, and then contact Isilon Technical Support.

500010005

SmartQuotas report generation failed.

Details

The Isilon SmartQuotas module failed to generate a requested quota report.

Administrator Response

Attempt to regenerate the quota report.

If you cannot regenerate the quota report, run the `isi_gather_info` command to gather logs, and then contact Isilon Technical Support.

Snapshot events

Snapshot events provide information about the snapshot daemon, such as status and configuration errors.

600010001

The snapshot daemon failed to create a scheduled snapshot.

Details

The system cannot create a requested snapshot. The request can be either a manual request or the result of a policy.

If the cluster is split, this message may appear on the minority group, the group that has fewer than half of the nodes. In this case, the error persists until the cluster is healthy, and you can safely ignore the error.

There are multiple ways to determine whether a node is in the majority group. One method is to run the `sysctl -a | grep has_quorum` command from the node that is reporting the error. If the command returns 1, the node is in the majority group.

Administrator Response

Complete the following steps:

1. Determine whether the number of snapshots exceeds the system-wide and directory limits. (The system-wide limit is 2048, and the directory limit is 1024.)
2. If the number of snapshots is at or exceeds the limits, either delete snapshots or contact Isilon Technical Services for assistance.
3. Confirm that snapshots on the directory do not exceed the quota.
4. If the issue is not resolved, run `isi_gather_info` to gather logs and contact Isilon Technical Support.

600010002

The snapshot daemon failed to delete an expired snapshot.

Details

The system failed to remove an existing snapshot.

If the cluster is split, this message may appear on the minority group, the group that has fewer than half of the nodes.

Administrator Response

Perform the following steps:

1. If the cluster is split, determine whether the error is occurring on the minority or the majority group.



Note: There are multiple ways to determine whether a node is in the majority group. One method is to run the `sysctl efs.gmp.has_quorum` command from the node that is reporting the error. If the command returns 1, the node is in the majority group.

If the error occurred on the minority group, the error is expected and may continue until the cluster is healthy. No further action is required.

2. If the error occurred on the majority group, perform the following tasks:
 - a. Verify that there is sufficient disk space on the cluster. If the cluster is more than 99% full, delete files or add space to the cluster.
 - b. Ensure that the restriper is running on all nodes. To determine whether the restriper is running on all nodes, log into any node via secure shell (SSH) connection or the serial console and run the following command:

```
isi_for_array -s 'pgrep isi_restripe_d|wc -l'|grep '^[0-9]0$'
```

Nodes that are listed in the output do not have the restriper running.

3. If the issue is not resolved, run the `isi_gather_info` to gather logs, and then contact Isilon Technical Support.

600010003

The snapshot daemon failed to remove a snapshot lock.

Details

The system cannot remove an expired snapshot lock. This error can occur when a disk is unwriteable.

If the cluster is split, this error may occur on the minority group, or the group that contains fewer than half of the nodes. In this case, the message persists until the cluster is healthy, and you can safely ignore the error.

In the unlikely event that this error occurs on the majority group, contact Isilon Technical Support for assistance.

Administrator Response

Perform the following steps:

1. If the cluster is split, determine whether the error is occurring on the minority or the majority group.

- If the error occurred on the majority group, contact Isilon Technical Support.
- If the error occurred on the minority group, you can safely ignore the error message.



Note: There are multiple ways to determine whether a node is in the majority group. One method is to run the `sysctl efs.gmp.has_quorum` command from the node that is reporting the error. If the command returns 1, the node is in the majority group.

2. If the cluster is not split, attempt the following possible solutions in the order presented. If a solution resolves the issue, there is no need to attempt other possible solutions:
 - Confirm that there is free disk space on the cluster. If the cluster is more than 99 percent full, delete files or add space to the cluster.
 - Ensure that the restriper is running on all nodes.

If neither of the possible solutions resolves the issue, contact Isilon Technical Support.

600010004

Snapshot daemon schedule policy config error.

Details

The `snapshot_schedule.xml` file is corrupt or unreadable.

Administrator Response

Run the `isi_gather_info` command to gather logs, and then contact Isilon Technical Support.

600010005

The snapshot reserve space is nearly full.

Details

The amount of data stored on the cluster is approaching or has exceeded the snapshot reserve space.

Exceeding the snapshot reserve space does not result in a failure to write snapshots to the cluster. The system can use any available disk space to write snapshots, and snapshots can exceed the snapshot reserve space.

However, problems occur when the available space in the cluster is less than the snapshot reserve space. If the cluster exceeds its snapshot reserve space, all attempts to write non-snapshot data to the cluster fail.

Administrator Response

Review cluster and snapshot utilization data, and then perform either of the following tasks:

- Delete some snapshots to reduce the amount of snapshot reserve space in use.
- Disable the snapshot reserve space. On Isilon clusters, snapshot reserve space is not required to write snapshots to disk.

Windows Networking events

Windows Networking events provide information about network services, such as Active Directory availability and networking component issues.

700010001

Windows ADS time skew exceeds threshold minutes. Authentication may fail.

Details

The time indicated on one or more nodes differs from the time indicated on the Active Directory clock by at least five minutes.

Administrator Response

Complete the following steps:

1. Check the time on the Active Directory server, and then perform one of the following tasks:
 - If the Active Directory time is incorrect, adjust the time on the Active Directory server.
 - If the Active Directory time is correct, adjust the cluster's time to match the time on the Active Directory server via the Isilon web administration interface.
2. If Network Time Protocol (NTP) is enabled, disable it. You cannot use both an Active Directory server and NTP to synchronize time on the same cluster.

If the issue is not resolved, run the `isi_gather_info` command to gather logs, and then contact Isilon Technical Support.

700010003

Windows time server could not be contacted.

Details

Windows time server could not be contacted.

Administrator Response

No action is required. The message is informational only.

700010004

SMB configuration upgrade failure

Details

SMB configuration upgrade failure.

Administrator Response

The message is informational only.

700020001

Windows UID map range is full.

Details

The user ID (UID) range for mapping Active Directory groups has run out of IDs and must be expanded.

Administrator Response

Contact Isilon Technical Support.

700020002

Windows GID map range is full. Authentication may fail until the range is increased.

Details

The group ID (GID) range for mapping Microsoft Active Directory groups has run out of IDs and must be expanded.

Administrator Response

Contact Isilon Technical Support.

700020003

Failed to parse idmap rules

Details

Failed to parse idmap rules.

Administrator Response

The message is informational only.

700030001

AD machine account missing.

Details

The Microsoft Active Directory account data that is stored on the cluster has been deleted or damaged.

Administrator Response

Run the following command:

```
isi auth ads accounts list
```

- If the output indicates that you must rejoin the domain, attempt to rejoin the cluster to the Active Directory Services domain. If you cannot do so, contact Isilon Technical Support.
- If the `isi auth ads accounts list` command produces any other output, contact Isilon Technical Support.

700030002

The Active Directory domain is offline. Authentication services may be interrupted.

Details

The node cannot contact an authentication server for the specified Active Directory domain.

The text of the event indicates the specific node on which the issue occurred.

By design, the node periodically attempts to re-establish communication with the domain. If the node successfully connects to the domain, the event clears itself.

Administrator Response

If the event does not clear itself within five minutes or if it recurs, perform the following steps on the node on which the issue occurred:

1. Attempt to ping the authentication server.
2. If the ping operation fails, confirm that network connectivity is active between the cluster and the authentication server, and then re-attempt the ping operation.

If you cannot restore network connectivity between the cluster and the authentication server to complete a successful ping operation, run the `isi_gather_info` command, and then contact Isilon Technical Support for assistance with troubleshooting.

3. If the ping operation succeeds, run the following command:

```
wbinfo -t
```

This command forces the node to attempt to establish communication with the domain.

4. If the `wbinfo -t` command output reports a failure, attempt to restart the authentication daemon by running the following command:

```
killall -TERM lsassd
```

- If the authentication daemon successfully restarts, the event clears itself. Proceed to step 5.
- If the authentication daemon does not successfully restart, run the `isi_gather_info` command, and then contact Isilon Technical Support.

5. If the `wbinfo -t` command reports success, attempt to map a drive from a client to the node on which the error occurred.

This step tests for end-to-end network connectivity, and the result indicates whether the error is transient.

- If you can map the drive without entering your username and password, the error is transient and there is no need to contact Isilon Technical Support.

However, a transient error may indicate an issue with Active Directory. Consider contacting your Active Directory administrator to determine whether any known issues may have caused the OneFS operating system to generate this event. Resolve issues with the Active Directory administrator.

- If the error is not transient or if it recurs, run the `isi_gather_info` command, and then contact Isilon Technical Support.

700030003

Authentication Provider initialization failure.

Details

The node cannot read from nor write to the Likewise authentication database files.

This event typically appears when a node does not have quorum or otherwise cannot access the Likewise authentication database files.

This event also sometimes appears when a node starts up. In that case, the event frequently resolves itself within five minutes. If the event resolves itself, no action is required.

Administrator Response

If the event does not resolve itself within five minutes, review the node's status via the Isilon web administration interface and resolve any apparent issues.

If the event still does not resolve itself, contact Isilon Technical Support.

700030004

Authentication service unavailable.

Details

The node cannot contact an authentication server for the specified domain.

The text of the event indicates the specific node on which the issue occurred.

By design, the node periodically attempts to re-establish communication with the domain. If the node successfully connects to the domain, the event clears itself.

Administrator Response

If the event does not clear itself within five minutes or if it recurs, perform the following steps on the node on which the issue occurred:

1. Attempt to ping the authentication server.
2. If the ping operation fails, confirm that network connectivity is active between the cluster and the authentication server, and then reattempt the ping operation.

If you cannot restore network connectivity between the cluster and the authentication server to complete a successful ping operation, run the `isi_gather_info` command, and then contact Isilon Technical Support for assistance with troubleshooting.

3. If the ping operation succeeds, run the following command:

```
wbinfo -t
```

This command forces the node to attempt to establish communication with the domain.

4. If the `wbinfo -t` command output reports a failure, attempt to restart the authentication daemon by running the following command:

```
killall -TERM lsassd
```

- If the authentication daemon successfully restarts, the event clears itself. Proceed to step 5.
- If the authentication daemon does not successfully restart, run the `isi_gather_info` command, and then contact Isilon Technical Support.

5. If the `wbinfo -t` command reports success, attempt to map a drive from a client to the node on which the error occurred.

This step tests for end-to-end network connectivity, and the result indicates whether the error is transient.

- If you can map the drive without entering your username and password, the error is transient and there is no need to contact Isilon Technical Support.

However, a transient error may indicate an issue with Active Directory. Consider contacting your Active Directory administrator to determine whether any known issues may have caused the OneFS operating system to generate this event. Resolve issues with the Active Directory administrator.

- If the error is not transient or if it recurs, run the `isi_gather_info` command, and then contact Isilon Technical Support .

700030005

AD server missing needed SPN(s) try 'isi auth ads spn check

Details

AD server missing needed SPN(s) try 'isi auth ads spn check.

Administrator Response

The message is informational only.

700030006

AD machine account invalid

Details

AD machine account invalid.

Administrator Response

The message is informational only.

700100001

Lwio Parameter Invalid

Details

Lwio Parameter Invalid.

Administrator Response

The message is informational only.

799910001

There are problems with Windows networking components.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

799920001

There are problems with the Windows ID map.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

799930001

There are problems with Windows authentication components.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

File System events

File System events provide information about file system errors.

800010002

DS fault detected.

Details

The system detected a metadata referential integrity error. Please contact support.

Administrator Response

Contact Isilon Technical Support.

800010003

IDI error in block.

Details

The system cannot verify data integrity.

Administrator Response

Contact Isilon Technical Support immediately.

800010004

IDI error: Shallow verification failure in block.

Details

The system cannot verify data integrity.

Administrator Response

Contact Isilon Technical Support immediately.

800010005

Repair attempt failed for file.

Details

A Dynamic Sector Repair (DSR) failed to reconstruct a block from parity data.

Administrator Response

Contact Isilon Technical Support.

800010006

System is running out of file descriptors.

Details

System is running out of file descriptors.

Administrator Response

No action is required. The message is informational only.

899990001

Filesystem problems detected.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

Hardware events

Hardware events provide information about hardware-specific status, such as voltage, power supply, and fan speed issues.

900010001

Node has had a clock failure.

Details

There is an error on the node motherboard, such as a faulty clock battery.

Administrator Response

Contact Isilon Technical Support.

900010002

Node has an NVRAM battery problem.

Details

The nonvolatile random access memory (NVRAM) voltage on a node is at an unexpected level.

Administrator Response

Contact Isilon Technical Support for help with troubleshooting.

900010003

Node has an NVRAM error.

Details

A process on the system attempted to read from the nonvolatile random access memory (NVRAM) board, but failed. The attempt resulted in an error-correcting code (ECC) error, which was corrected.

Administrator Response

Contact Isilon Technical Support.

900010004

Node has an open chassis.

Details

A chassis sensor indicates that a node is open.

This alert normally appears when maintenance is being performed and a node is open. While the node is open, it is set to read-only status. When maintenance is complete, the alert clears itself.

Administrator Response

Confirm whether maintenance is being performed on the node. If it is, and if the alert persists after maintenance is complete, use the Isilon web administration interface to clear the alert manually.

If the alert does not clear itself when maintenance is complete, or if maintenance is not being performed on the node, contact Isilon Technical Support.

900010005

DMI log entry

Details

A node indicates a memory or Peripheral Component Interconnect (PCI) bus error.

Administrator Response

Contact Isilon Technical Support.

900010006

DMI read failure.

Details

A node indicates a memory or Peripheral Component Interconnect (PCI) bus error.

Administrator Response

Contact Isilon Technical Support.

900010007

Policy Error -- DIMM, Correctable ECC error policy violation.

Details

A memory module in the node has exceeded Isilon's policy for recoverable failures and should be replaced.

Administrator Response

Schedule a replacement of the Dual Inline Memory Module (DIMM) with Isilon Technical Support.

900010008

System management hardware failure.

Details

The I2C bus hardware on the node is having problems.

Administrator Response

Use the Isilon web administration interface to clear the alert manually.

If clearing the alert does not resolve the issue, contact Isilon Technical Support.

900020001

Hardware Monitoring failure.

Details

The subsystem that monitors the health of the hardware (such as the temperature and fan speeds) has failed.

This event can occur intermittently without harm to the system.

Administrator Response

If the event occurs intermittently, you can safely ignore it.

If the event persists or recurs frequently, contact Isilon Technical Support.

900020026

CPU 0 about to throttle due to temperature.

Details

A temperature sensor on a node indicates an elevated temperature.

Administrator Response

Attempt the following possible solutions. If a solution resolves the issue, there is no need to attempt other possible solutions:

- Review the chassis inlet temperature statistics, which are included in the event. If the temperature is elevated, the problem is likely a high ambient temperature in the data center. Address any environmental concerns.
- Run the following command:

```
isi_hw_status
```

Review the output to determine whether there is a slow or failed fan that was not otherwise reported. To order a replacement fan, contact Isilon Technical Support.

- Check the temperature of the data center and the airflow within the rack. Address any environmental concerns.
- Check for high CPU and disk utilization in the node. High utilization can contribute to high temperatures within the node.

Cancel or quiet the event.

900020027

CPU 1 about to throttle due to temperature.

Details

A temperature sensor on a node indicates an elevated temperature.

Administrator Response

Attempt the following possible solutions. If a solution resolves the issue, there is no need to attempt other possible solutions:

- Review the chassis inlet temperature statistics, which are included in the event. If the temperature is elevated, the problem is likely a high ambient temperature in the data center. Address any environmental concerns.
- Run the following command:

`isi_hw_status`

Review the output to determine whether there is a slow or failed fan that was not otherwise reported. To order a replacement fan, contact Isilon Technical Support.

- Check the temperature of the data center and the airflow within the rack. Address any environmental concerns.
- Check for high CPU and disk utilization in the node. High utilization can contribute to high temperatures within the node.

Cancel or quiet the event.

900020033

Redundant power supply failure.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.


The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)


- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.

 **Caution:** Do not move the power cable to another power supply for the same node.
 - c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.

 **Caution:** Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.

 - If the issue follows the power supply, the power supply needs to be replaced.
 - If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.
 - d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.
2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900020034

Physical Memory low.

Details

A node reports less than the expected amount of physical memory.

Administrator Response

Contact Isilon Technical Support.

900020035

CPU throttling.

Details

A temperature sensor on a node indicates an elevated temperature.

Administrator Response

Attempt the following possible solutions. If a solution resolves the issue, there is no need to attempt other possible solutions:

- Review the chassis inlet temperature statistics, which are included in the event. If the temperature is elevated, the problem is likely a high ambient temperature in the data center. Address any environmental concerns.
- Run the following command:

```
isi_hw_status
```

Review the output to determine whether there is a slow or failed fan that was not otherwise reported. To order a replacement fan, contact Isilon Technical Support.

- Check the temperature of the data center and the airflow within the rack. Address any environmental concerns.
- Check for high CPU and disk utilization in the node. High utilization can contribute to high temperatures within the node.

Cancel or quiet the event.

900030001

Hardware Monitoring failure.

Details

The subsystem that monitors the health of the hardware (such as the temperature and fan speeds) has failed.

This event can occur intermittently without harm to the system.

Administrator Response

If the event occurs intermittently, you can safely ignore it.

If the event persists or recurs frequently, contact Isilon Technical Support.

900030022

Physical Memory low.

Details

A node reports less than the expected amount of physical memory.

Administrator Response

Contact Isilon Technical Support.

900030023

CPU throttling.

Details

A temperature sensor on a node indicates an elevated temperature.

Administrator Response

Attempt the following possible solutions. If a solution resolves the issue, there is no need to attempt other possible solutions:

- Review the chassis inlet temperature statistics, which are included in the event. If the temperature is elevated, the problem is likely a high ambient temperature in the data center. Address any environmental concerns.
- Run the following command:

```
isi_hw_status
```

Review the output to determine whether there is a slow or failed fan that was not otherwise reported. To order a replacement fan, contact Isilon Technical Support.

- Check the temperature of the data center and the airflow within the rack. Address any environmental concerns.
- Check for high CPU and disk utilization in the node. High utilization can contribute to high temperatures within the node.

Cancel or quiet the event.

900040001

Hardware Monitoring failure.

Details

The subsystem that monitors the health of the hardware (such as the temperature and fan speeds) has failed.

This event can occur intermittently without harm to the system.

Administrator Response

If the event occurs intermittently, you can safely ignore it.

If the event persists or recurs frequently, contact Isilon Technical Support.

900040034

Physical Memory low.

Details

A node reports less than the expected amount of physical memory.

Administrator Response

Contact Isilon Technical Support.

900040035

CPU throttling.

Details

A temperature sensor on a node indicates an elevated temperature.

Administrator Response

Attempt the following possible solutions. If a solution resolves the issue, there is no need to attempt other possible solutions:

- Review the chassis inlet temperature statistics, which are included in the event. If the temperature is elevated, the problem is likely a high ambient temperature in the data center. Address any environmental concerns.
- Run the following command:

```
isi_hw_status
```

Review the output to determine whether there is a slow or failed fan that was not otherwise reported. To order a replacement fan, contact Isilon Technical Support.

- Check the temperature of the data center and the airflow within the rack. Address any environmental concerns.
- Check for high CPU and disk utilization in the node. High utilization can contribute to high temperatures within the node.

Cancel or quiet the event.

900060001

Hardware Monitoring failure.

Details

The subsystem that monitors the health of the hardware (such as the temperature and fan speeds) has failed.

This event can occur intermittently without harm to the system.

Administrator Response

If the event occurs intermittently, you can safely ignore it.

If the event persists or recurs frequently, contact Isilon Technical Support.

900060026

CPU throttling.

Details

A temperature sensor on a node indicates an elevated temperature.

Administrator Response

Attempt the following possible solutions. If a solution resolves the issue, there is no need to attempt other possible solutions:

- Review the chassis inlet temperature statistics, which are included in the event. If the temperature is elevated, the problem is likely a high ambient temperature in the data center. Address any environmental concerns.
- Run the following command:

```
isi_hw_status
```

Review the output to determine whether there is a slow or failed fan that was not otherwise reported. To order a replacement fan, contact Isilon Technical Support.

- Check the temperature of the data center and the airflow within the rack. Address any environmental concerns.
- Check for high CPU and disk utilization in the node. High utilization can contribute to high temperatures within the node.

Cancel or quiet the event.

900060027

Redundant power supply failure.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.

The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)

- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:

- a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
- b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.



Caution: Do not move the power cable to another power supply for the same node.

- c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.



Caution: Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.

- If the issue follows the power supply, the power supply needs to be replaced.
- If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.

- d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.

2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:

- Supplied voltage.
- Power quality, such as the correctness of the voltage and the smoothness of the waves.
- Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
- UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900060028

Physical Memory low.

Details

A node reports less than the expected amount of physical memory.

Administrator Response

Contact Isilon Technical Support.

900060029

Power Supply 1 Over Current out of spec.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.



The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)

- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.

- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
 **Caution:** Do not move the power cable to another power supply for the same node.
 - c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.
 **Caution:** Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.
 - If the issue follows the power supply, the power supply needs to be replaced.
 - If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.
 - d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.
2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900060030

Power Supply 2 Over Current out of spec.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.


The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)


- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.

 **Caution:** Do not move the power cable to another power supply for the same node.
 - c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.

 **Caution:** Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.

 - If the issue follows the power supply, the power supply needs to be replaced.
 - If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.
 - d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.
2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900060031

Power Supply 1 Under Voltage.

Details

The issue is caused by one of the following conditions:



- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.

The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)

- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
 **Caution:** Do not move the power cable to another power supply for the same node.
 - c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.
 **Caution:** Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.
 - If the issue follows the power supply, the power supply needs to be replaced.
 - If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.
 - d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.
2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900060032

Power Supply 2 Under Voltage.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.

The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)

- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:

- a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
- b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.



Caution: Do not move the power cable to another power supply for the same node.

- c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.



Caution: Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.

- If the issue follows the power supply, the power supply needs to be replaced.
- If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.

- d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.

2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:

- Supplied voltage.
- Power quality, such as the correctness of the voltage and the smoothness of the waves.
- Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
- UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900060033

Power Supply 1 Over Voltage.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.

The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)

- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.

- b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.



Caution: Do not move the power cable to another power supply for the same node.

- c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.



Caution: Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.

- If the issue follows the power supply, the power supply needs to be replaced.
- If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.

- d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.

2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:

- Supplied voltage.
- Power quality, such as the correctness of the voltage and the smoothness of the waves.
- Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
- UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900060034

Power Supply 2 Over Voltage.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.

The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)

- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.



Caution: Do not move the power cable to another power supply for the same node.

- c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.



Caution: Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.

- If the issue follows the power supply, the power supply needs to be replaced.
 - If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.
- d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.
2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900060035

Power Supply 1 Fan Fail.

Details

The speed of one or more fans has fallen below an expected threshold.

Administrator Response

Contact Isilon Technical Support.

900060036

Power Supply 2 Fan Fail.

Details

The speed of one or more fans has fallen below an expected threshold.

Administrator Response

Contact Isilon Technical Support.

900060037

Power Supply 1 A/C Fail.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.



The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)

- Steady green (applies to all node models): Status is good.

- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
 **Caution:** Do not move the power cable to another power supply for the same node.
 - c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.
 **Caution:** Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.
 - If the issue follows the power supply, the power supply needs to be replaced.
 - If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.
 - d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.
2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900060038

Power Supply 2 A/C Fail.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.

The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)


- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.


- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.

 **Caution:** Do not move the power cable to another power supply for the same node.
 - c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.

 **Caution:** Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.

 - If the issue follows the power supply, the power supply needs to be replaced.
 - If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.
 - d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.
2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900080001

Hardware Monitoring failure.

Details

The subsystem that monitors the health of the hardware (such as the temperature and fan speeds) has failed.

This event can occur intermittently without harm to the system.

Administrator Response

If the event occurs intermittently, you can safely ignore it.

If the event persists or recurs frequently, contact Isilon Technical Support.

900080033

Redundant power supply failure.

Details

The issue is caused by one of the following conditions:



- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.

The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)

- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.
 **Caution:** Do not move the power cable to another power supply for the same node.
 - c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.
 **Caution:** Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.
 - If the issue follows the power supply, the power supply needs to be replaced.
 - If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.
 - d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.
2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900080034

Physical Memory low.

Details

A node reports less than the expected amount of physical memory.

Administrator Response

Contact Isilon Technical Support.

900080035

CPU throttling.

Details

A temperature sensor on a node indicates an elevated temperature.

Administrator Response

Attempt the following possible solutions. If a solution resolves the issue, there is no need to attempt other possible solutions:

- Review the chassis inlet temperature statistics, which are included in the event. If the temperature is elevated, the problem is likely a high ambient temperature in the data center. Address any environmental concerns.
- Run the following command:

```
isi_hw_status
```

Review the output to determine whether there is a slow or failed fan that was not otherwise reported. To order a replacement fan, contact Isilon Technical Support.

- Check the temperature of the data center and the airflow within the rack. Address any environmental concerns.
- Check for high CPU and disk utilization in the node. High utilization can contribute to high temperatures within the node.

Cancel or quiet the event.

900090001

Hardware Monitoring failure.

Details

The subsystem that monitors the health of the hardware (such as the temperature and fan speeds) has failed.

This event can occur intermittently without harm to the system.

Administrator Response

If the event occurs intermittently, you can safely ignore it.

If the event persists or recurs frequently, contact Isilon Technical Support.

900090023

Redundant power supply failure.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.

The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)

- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.



Caution: Do not move the power cable to another power supply for the same node.

- c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.



Caution: Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.

- If the issue follows the power supply, the power supply needs to be replaced.
 - If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.
- d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.
2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900090024

Physical Memory low.

Details

A node reports less than the expected amount of physical memory.

Administrator Response

Contact Isilon Technical Support.

900090025

CPU throttling.

Details

A temperature sensor on a node indicates an elevated temperature.

Administrator Response

Attempt the following possible solutions. If a solution resolves the issue, there is no need to attempt other possible solutions:

- Review the chassis inlet temperature statistics, which are included in the event. If the temperature is elevated, the problem is likely a high ambient temperature in the data center. Address any environmental concerns.
- Run the following command:

```
isi_hw_status
```

Review the output to determine whether there is a slow or failed fan that was not otherwise reported. To order a replacement fan, contact Isilon Technical Support.

- Check the temperature of the data center and the airflow within the rack. Address any environmental concerns.
- Check for high CPU and disk utilization in the node. High utilization can contribute to high temperatures within the node.

Cancel or quiet the event.

900090046

Redundant power supply failure on expansion chassis.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.

The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)

- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.



Caution: Do not move the power cable to another power supply for the same node.

- c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.



Caution: Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.

- If the issue follows the power supply, the power supply needs to be replaced.
 - If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.
- d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.
2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900110001

CPU throttling.

Details

A temperature sensor on a node indicates an elevated temperature.

Administrator Response

Attempt the following possible solutions. If a solution resolves the issue, there is no need to attempt other possible solutions:

- Review the chassis inlet temperature statistics, which are included in the event. If the temperature is elevated, the problem is likely a high ambient temperature in the data center. Address any environmental concerns.
- Run the following command:

`isi_hw_status`

Review the output to determine whether there is a slow or failed fan that was not otherwise reported. To order a replacement fan, contact Isilon Technical Support.

- Check the temperature of the data center and the airflow within the rack. Address any environmental concerns.
- Check for high CPU and disk utilization in the node. High utilization can contribute to high temperatures within the node.

Cancel or quiet the event.

900110002

Hardware Monitoring failure.

Details

The subsystem that monitors the health of the hardware (such as the temperature and fan speeds) has failed.

This event can occur intermittently without harm to the system.

Administrator Response

If the event occurs intermittently, you can safely ignore it.

If the event persists or recurs frequently, contact Isilon Technical Support.

900110003

Physical Memory low.

Details

A node reports less than the expected amount of physical memory.

Administrator Response

Contact Isilon Technical Support.

900110004

Power supply 1 failure.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.

The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)

- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.



Caution: Do not move the power cable to another power supply for the same node.

- c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.



Caution: Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.

- If the issue follows the power supply, the power supply needs to be replaced.
- If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.

- d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.
2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900110005

Power supply 2 failure.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.


The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)

- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.


Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.



Caution: Do not move the power cable to another power supply for the same node.
 - c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.



Caution: Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.
 - If the issue follows the power supply, the power supply needs to be replaced.
 - If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.
 - d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.

2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900120001

CPU throttling.

Details

A temperature sensor on a node indicates an elevated temperature.

Administrator Response

Attempt the following possible solutions. If a solution resolves the issue, there is no need to attempt other possible solutions:

- Review the chassis inlet temperature statistics, which are included in the event. If the temperature is elevated, the problem is likely a high ambient temperature in the data center. Address any environmental concerns.
- Run the following command:

```
isi_hw_status
```

Review the output to determine whether there is a slow or failed fan that was not otherwise reported. To order a replacement fan, contact Isilon Technical Support.

- Check the temperature of the data center and the airflow within the rack. Address any environmental concerns.
- Check for high CPU and disk utilization in the node. High utilization can contribute to high temperatures within the node.

Cancel or quiet the event.

900120002

Hardware Monitoring failure.

Details

The subsystem that monitors the health of the hardware (such as the temperature and fan speeds) has failed.

This event can occur intermittently without harm to the system.

Administrator Response

If the event occurs intermittently, you can safely ignore it.

If the event persists or recurs frequently, contact Isilon Technical Support.

900120003

Physical Memory low.

Details

A node reports less than the expected amount of physical memory.

Administrator Response

Contact Isilon Technical Support.

900120004

Power supply 1 failure.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.

The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)

- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.



Caution: Do not move the power cable to another power supply for the same node.

- c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.



Caution: Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.

- If the issue follows the power supply, the power supply needs to be replaced.
 - If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.
- d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.
2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

900120005

Power supply 2 failure.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.


The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)


- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.

 **Caution:** Do not move the power cable to another power supply for the same node.
 - c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.

 **Caution:** Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.

 - If the issue follows the power supply, the power supply needs to be replaced.
 - If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.
 - d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.
2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

910100001

Sensor out of spec.

Details

The speed of one or more fans has fallen below an expected threshold.

Administrator Response

Contact Isilon Technical Support.

910100002

Sensor out of spec.

Details

One or more sensors reports unexpected voltage levels.

Administrator Response

Contact Isilon Technical Support.

910100003

Sensor out of spec.

Details

A temperature sensor on a node indicates an elevated temperature.

Administrator Response

Attempt the following possible solutions. If a solution resolves the issue, there is no need to attempt other possible solutions:

- Review the chassis inlet temperature statistics, which are included in the event. If the temperature is elevated, the problem is likely a high ambient temperature in the data center. Address any environmental concerns.
- Run the following command:

`isi_hw_status`

Review the output to determine whether there is a slow or failed fan that was not otherwise reported. To order a replacement fan, contact Isilon Technical Support.

- Check the temperature of the data center and the airflow within the rack. Address any environmental concerns.
- Check for high CPU and disk utilization in the node. High utilization can contribute to high temperatures within the node.

Cancel or quiet the event.

910100004

Sensor out of spec.

Details

The issue is caused by one of the following conditions:

- One of the power supplies in the node is not providing power.
- The power (AC) provided from the electrical outlet to one or both power supplies is insufficient.


The LED light indicates the power supply status. (The LED light is on the inside of the node, but is visible on the back panel.)


- Steady green (applies to all node models): Status is good.
- Blinking green (applies only to 4U (36000X, 3600NL, 72000X, and 72000NL) nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Steady amber (applies only to all X- and S-series nodes): The electrical outlet is providing sufficient AC to the power supplies, but the node is turned off.
- Blinking amber (applies only to all X- and S-series nodes): There is a power supply failure.
- Light is off (applies to all node models): There is no or insufficient AC power.

Administrator Response

Perform the following steps:

1. If a single node reports the issue, determine the cause of the issue by performing the following steps:
 - a. Determine if the electrical outlet is functioning properly by plugging the power cable into a different electrical outlet.
 - b. If the issue is not resolved by using a different electrical outlet, move the power cable from the power supply that reports the failure to the power supply of a node that does not report a failure. If the cable is the issue, replace the cable.

 **Caution:** Do not move the power cable to another power supply for the same node.
 - c. If the issue persists, take one power supply out of a different node and replace it with the power supply that reports the problem.

 **Caution:** Do not switch power supplies in the same node. Doing so causes there to be a time when neither power supply is in the node.

 - If the issue follows the power supply, the power supply needs to be replaced.
 - If the issue stays with the node, contact Isilon Technical Support for help troubleshooting the node.
 - d. If the preceding steps do not resolve the issue, contact Isilon Technical Support.
2. If multiple nodes report power supply issues, confirm the health of power subsystem. The issue is likely environmental. Check each of the following items:
 - Supplied voltage.
 - Power quality, such as the correctness of the voltage and the smoothness of the waves.
 - Uninterruptible Power Supply (UPS) health, such as the ability to provide clean and consistent power.
 - UPS battery health.

If you need help troubleshooting the issue, contact Isilon Technical Support.

999910001

There are hardware issues.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.

3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

999910002

There are sensor issues.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

999910003

There are fan issues.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

999910004

There are temperature issues.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.

999910005

There are voltage issues.

Details

A series of related events has occurred.

Administrator Response

Perform the following steps:

1. In the **Actions** column for the event, select **View**.
2. The **Events > View Event** page appears.
3. In the **Coalesced Events** section, review the individual events related to this event. If you cannot identify the problem, contact Isilon customer support.